

# Strong Post-Compromise Secure Proxy Re-Encryption

Alex Davidson<sup>†</sup>, Amit Deo<sup>†</sup>, Ela Lee<sup>†,1</sup>, Keith Martin

Royal Holloway, University of London

**Abstract** Proxy Re-Encryption (PRE), introduced by Bellare et. al in [6], allows a ciphertext encrypted using a key  $pk_i$  to be re-encrypted by a third party so that it is an encryption of the same message under a new key  $pk_j$ , without revealing the message. Post-Compromise Security (PCS) was first introduced for messaging protocols, and ensures that a ciphertext remains confidential even when past keys have been corrupted. We define PCS in the context of PRE, which ensures that an adversary cannot distinguish which ciphertext a re-encryption was created from even given the old secret key, potential old ciphertexts and update token used to perform the re-encryption. We argue that this formal notion accurately captures the most intuitive form of PCS. We give separating examples demonstrating how our definition is stronger than existing ones, before showing that PCS can be met using a combination of existing security definitions from the literature. In doing so, we show that there are existing PRE schemes that satisfy PCS. We also show that natural modifications of more practical PRE schemes can be shown to be PCS without relying on this combination of existing security definitions. Finally, we discuss the relationship between PCS with selective versus adaptive key corruptions, giving a theorem that shows how adaptive security can be met for certain re-encryption graphs.

## 1 Introduction

Cloud storage has become increasingly popular in recent years. Not long ago, cloud storage was mainly viewed as a useful backup for any files stored locally. However, in some situations cloud storage has started to become the for all storage and not just as a backup. One of the most prolific examples of this is in the increased popularity of media streaming platforms such as Netflix and Spotify, where clients subscribe to a service that gives them access to files on demand as opposed to storing such files locally. Other advantages come in the surge of smart devices, which no longer need much in-built storage as long as they have an internet connection.

---

<sup>†</sup> These authors are supported by the EPSRC and the UK government as part of the Centre for Doctoral Training in Cyber Security at Royal Holloway, University of London (EP/K035584/1).

<sup>1</sup> Contact author: Ela.Lee.2010@live.rhul.ac.uk

Since the cloud is a third party, clients must encrypt their files to control ownership and access to their data. The difficulty comes when the key which a file is encrypted under needs changing. Key changes can happen for a variety of reasons. For instance: as a matter of good practice; or to enforce access control. For the former, NIST recommends regular key rotation [3], as does the Payment Card Industry Data Security Standard [29] and the Open Web Application Security Project (OWASP) [28]. For the latter, an organisation can choose which of its employees can access specific files by having those files encrypted under those keys. Should access need to be granted or revoked from someone, the files should be re-encrypted to a new key accordingly.

One trivial solution is to have the client download, decrypt, encrypt using the new key, and then re-upload the file. However, this can be very expensive in terms of both computational resources and storage, particularly for modern applications involving large databases, or if the client has limited processing capability. The primitive of Proxy Re-Encryption (PRE), introduced by Blaze et al. [6], presents a more elegant solution. In a PRE scheme, client creates an *update token*  $\Delta_{i,j}$  using the current secret key  $sk_i$  and new public key  $pk_j$  and sends it to the server. The server can then use this token to re-encrypt the ciphertext, transforming it into an encryption of the same message which can now be decrypted using  $sk_j$ . The most basic security notion for PRE states that the server performing the re-encryption learns nothing about the underlying message.

**Existing security definitions.** The most basic security definition for PRE is that the server should not learn the underlying message or either of the keys – essentially an extension of Indistinguishable against Chosen Plaintext Attacks (IND-CPA) security factoring in re-encryption capabilities. Much existing work focused on creating definitions for CPA and Chosen Ciphertext Attacks (CCA) security for PRE variants, considering unidirectional/bidirectional PRE and single/multi-hop PRE. Such work rarely considers revocation as an application.

There is little work on requiring re-encryption to also re-randomise. [22] has one such definition IND-UPD for *updatable encryption schemes*, which are similar to symmetric PRE schemes but assume keys are updated sequentially. [16] also has similar definition UP RE-ENC. However, these definitions are limited to the key rotation application, and consider some schemes to be secure which we believe do not believe meet the intuition of PCS. We give further detail in Section 2.

**Post-Compromise Security (PCS).** The notion of PCS was first used in [13] for messaging protocols, which is informally defined as follows:

**Definition 1 ([13]).** *A protocol between Alice and Bob provides Post-Compromise Security (PCS) if Alice has a security guarantee about communication with Bob, even if Bob’s secrets have already been compromised.*

As this definition gives the bare intuition, the ‘security guarantee’ is left open so that it can be tailored to the application, as we do in this work. This differs from *forward security*, which conveys that the compromise of future states does not affect the security of past ones. PCS conveys a scheme’s ability to regain security after a session or party has been compromised, which has clear applications

involving revocation and key rotation (key life-cycles), where having access to the old key should not affect the security of a re-encrypted ciphertext.

**Motivation for PCS PRE.** One particular application of post-compromise PRE of interest is to complement PCS of messages in transit, by giving PCS security to backed up messages stored in the cloud. The Signal Protocol for encrypting messages in transit between two devices provides strong, modern, provable security guarantees including PCS [13]. However, most users expect to keep their messages when they lose their device or buy a new one; thus, popular Signal implementations such as WhatsApp back up client messages to public cloud services. Unfortunately, this backup is encrypted using a static encryption key. This means that while messages in transit have PCS, these properties are lost once messages are backed up. If an adversary compromises a device and obtains the static cloud backup key, they can retain this to compromise future messages once they are backed up.

Assuming that all message history is stored locally, the updated message history could be encrypted under a new key and re-uploaded at regular time intervals, but this will have a huge cost both in terms of computation and bandwidth, particularly as much messaging is done via smart-phones. A PRE scheme with PCS could be used instead, so that the PCS of messages in transit is extended to message backups.

## 1.1 Contributions

In this paper we set out the first formalisation of PCS for PRE schemes. Considering PCS for PRE, a strong definition would consider what happens when both old keys and update tokens used to perform a re-encryption are compromised. However, to date there is no security definition for that does this. We create a new definition which implies both unidirectionality, and that additional randomness beyond that given in the update token is added upon re-encryption. In our model, the adversary cannot distinguish a re-encrypted ciphertext given the old key, old ciphertexts and the token used to perform the re-encryption. In other words, we view a compromise as the loss of all previous public and secret states associated with a given ciphertext. Overall, we limit the information that must remain secret to the current secret key alone. Since we do not make as many assumptions on which algorithms are deterministic or on the flow of re-encryption operations, our definition applies to more general PRE schemes than similar definitions in the literature (see Section 2).

We analyse our model, proving several results that associate PCS with existing security models for PRE, as well as providing separating examples that distinguish PCS as a separate security characteristic in its own right. One of our major contributions is to show that a PRE scheme that is both *source-hiding* and secure against *honest re-encryption attacks* also has PCS, meaning that several PRE schemes given by Fuchsbauer et al. [17] immediately satisfy PCS.

Unfortunately, these schemes require inefficient parameter choices, and are only proven secure for restricted re-encryption graphs (the structure by which

ciphertexts are moved between keys, see ??) with a single source node, since the work of [17] focuses on achieving different security goals. We therefore give a new construction of a PRE scheme, pcBV-PRE, adapted from BV-PRE – the performant ring-Learning With Errors (RLWE) scheme of Polyakov et al. [30]. We prove that our adaptation achieves PCS and IND-CPA-security independently of the transformation that we define above, meaning that source-hiding is not necessary to achieve PCS. The advantage of this is that we can prove security for more general re-encryption graphs. Moreover, the lattice-based source-hiding schemes of [17] are forced into sub-optimal parameter choices that render their assumptions much stronger (with respect to the approximation factors of solving worst-case lattice problems) and much less efficient. Our new scheme is much more efficient since we leverage the speed of the original construction with minimal changes.

Finally, we show that achieving PCS with adaptive key compromises is possible via the transformation of [17] using adaptively IND-CPA-secure and source-hiding PRE schemes. Such schemes currently only exist for certain graphs (trees and chains) with polynomial security loss. We leave achieving adaptively-secure PCS for more general graphs as interesting future work.

**Paper structure.** We begin by reviewing related work and necessary preliminaries in Sections 2–4. Motivated by our discussion of related work, we define a new notion of PCS for PRE in Section 5 before demonstrating how our definition relates to those already in the literature using separating examples to demonstrate that our definition is stronger than previous work, then showing how a combination of existing definitions leads to PCS. In Section 6, we give an explicit, efficient construction by modifying the lattice-based BV-PRE scheme [30] to show that our notion of PCS can be satisfied by natural extensions of current practical PRE schemes. Finally, in Section 7 we discuss the relationship between selective and adaptive security for PCS using the work of [17].

## 2 Related work

There are many different types of PRE schemes in the literature that satisfy various properties. For example, PRE schemes can be bidirectional, meaning that a single update token for Alice and Bob allows to re-encrypt ciphertexts not only from Alice to Bob, but from Bob to Alice too. Alternatively, PRE schemes can be unidirectional in the sense that an update token only allows one to update ciphertexts in one direction. There is also a distinction between single-hop and multi-hop PRE schemes. Single-hop schemes only allow ciphertexts to be re-encrypted once whereas multi-hop schemes allow for multiple re-encryptions.

The basic security definition for PRE was first given in [6] for bidirectional schemes. It is similar to IND-CPA security but adapted to allow for re-encryption capabilities, meaning that the proxy does not learn the underlying message during the re-encryption process. In this work we refer to such notions as PRE-CPA, to avoid confusion with IND-CPA (for Public Key Encryption (PKE) schemes) and PKE-CPA (for PRE schemes). Unidirectional PRE schemes were introduced

by Ateniese et al. [2] together with an equivalent security notion. A definition for Indistinguishable against Chosen Ciphertext Attacks (IND-CCA) security for PRE first appears in [9] for bidirectional single-hop (ciphertexts can only be re-encrypted once) schemes, and is the first definition where the adversary can adaptively corrupt secret keys. A definition for CCA security for unidirectional schemes is given in [23].

It was an open problem for many years to create a PRE scheme which is both unidirectional and multi-hop, with single-hop constructions emerging as the means of providing unidirectionality in the meantime. Multi-hop schemes are necessary for both key rotation and dynamic access control. Unidirectional schemes are necessary for applications where the trust relationship between Alice and Bob is not symmetrical, such as if Bob is more senior to Alice and therefore has access to more sensitive files. Whilst most existing literature considers directionality as a class of PRE schemes, it was been defined explicitly as a formal security definition in [21]. The first constructions for PRE-CPA secure unidirectional, multi-hop PRE schemes were given using program obfuscation [11, 19], until Gentry [18] gave a generic construction using Fully Homomorphic Encryption (FHE).

A significant recent development in recent times is the notion of Honest Re-encryption Attacks (HRA), introduced by Cohen [12] which allows the adversary to re-encrypt non-challenge ciphertexts from uncorrupted keys to corrupted ones, unlike in PRE-Indistinguishable against Chosen Plaintext Attacks (PRE-CPA) where no such re-encryptions are permitted. Cohen also defines *re-encryption simulatability* and demonstrates IND-CPA-secure schemes which have this property are Indistinguishable against Honest Re-encryption Attacks (IND-HRA)-secure.

Fuchsbauer et al. present relations between selective and adaptive security in [17]. They expand on Cohen’s work, defining *weak key privacy* and *source-hiding* and show how these can reduce security against adaptive HRAs to security against CPA, with sub-exponential loss for some re-encryption graphs (see Section 3 for a definition of re-encryption graphs). In particular, they show quasi-polynomial loss in security when lifting to adaptive corruption for trees and chains, and exponential loss otherwise. However it is also worth noting that their proofs rely on re-encryption trees having only one source, which limits the number of applications the scheme is guaranteed to be secure in.

The original notion of post-compromise security is from messaging protocols [13], to reflect the idea that compromise of past sessions does not impact the confidentiality of future ones. PCS builds on the idea that re-encryption should fully re-randomise a ciphertext upon re-encryption. To see this, consider the *key encapsulation approach* to PRE: encryption involves generating a symmetric key  $k$  and performing a symmetric encryption  $C_1 \stackrel{\$}{\leftarrow} \mathcal{SE}.Enc(k, m)$ , before encrypting  $k$  using a public-key PRE scheme  $C_0 \leftarrow \mathcal{PRE}.Enc(pk, k)$ . For re-encryption, the ciphertext header  $C_0$  is re-encrypted so it is now an encryption of  $k$  under some specified  $pk'$ , whilst  $C_1$  remains unchanged. This method is popular as it allows for a hybrid approach, where encryption of the message is symmetric and the header uses public-key encryption. With this approach, key-scraping attacks are possible. A malicious user simply retains the symmetric encryption key  $k$  and

can derive the message regardless of how often the ciphertext is re-encrypted. It may be unrealistic for a malicious revoked user to download all the plaintexts due to storage constraints, as is the case for subscriber-based streaming platforms. However, as symmetric keys are typically much shorter than the plaintexts, it is much more realistic that a malicious subscriber could perform the above attack.

Whilst re-randomisation is not generally a goal of PRE schemes, it has been covered recently in related work [9, 16, 22]. [9] discusses a notion for re-randomising ciphertexts (which they call *unlinkability*) only in the context of creating a CCA definition which considers re-randomisation. Other notions of re-randomisation mainly exist in the related areas of *key rotation* and *updatable encryption*. These differ from proxy re-encryption in that they are symmetric schemes, that re-encryption happens sequentially from  $k_i$  to  $k_{i+1}$ , and they name full re-randomisation as a necessary security goal.

Definitions conveying this include IND-UPD security [22] and UP-REENC security [16]. IND-UPD aims to cover PCS for *updatable encryption* schemes. However, an updatable encryption scheme can be bidirectional and still be IND-UPD, as bidirectionality invokes the additional winning condition that the adversary cannot have learned any update tokens used to create a challenge ciphertext. We give an asymmetric version of IND-UPD in Appendix B, which we use in Section 5.4 to demonstrate that our notion is stronger. UP-REENC security does not give the adversary the token used to create the challenge re-encryption, and therefore does not cover full compromise of the user who generated the update token. Their construction, ReCrypt allows an adversary who has learned the update token to reverse the re-encryption of the challenge ciphertext. Other related work models PCS by giving a bound on the amount of information the adversary maintains of the old ciphertext [21, 27]. However, such schemes are weaker than the notion we define particularly if considering revoked users storing parts of the original ciphertext and colluding, or a user relinquishing ownership of a file. We create a stronger definition in which schemes must be unidirectional, and the adversary knows the update token used to create the challenge re-encryption.

### 3 Preliminaries

In this section, we give the preliminaries for proxy re-encryption, including some common security definitions and an explanation of directed re-encryption graphs. Whilst we stick to the asymmetric setting in the body of this work, we give symmetric variants for important definitions in Appendix A for easier comparison with related work in the symmetric setting.

**Definition 2.** A Proxy Re-Encryption (PRE) scheme consists of the following algorithms:

- $\text{Setup}(1^\lambda) \rightarrow \text{params}$ : Outputs a set of public parameters, including the message space and ciphertext space. Note that  $\text{params}$  is input to every subsequent algorithm, but we leave it out for compactness of notation. We often omit the Setup algorithm for the same reason.

- $\text{KeyGen}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$ : Generates a public-private key pair.
- $\text{Enc}(\text{pk}, m) \rightarrow C$ : Encrypts a message  $m$  using a public key  $\text{pk}$ , producing a ciphertext  $C$ .
- $\text{Dec}(\text{sk}, C) \rightarrow m' \cup \perp$ : Decrypts a ciphertext  $C$  to produce either an element of the message space  $m'$  or an error symbol  $\perp$ .
- $\text{ReKeyGen}(\text{sk}_i, \text{pk}_j) \rightarrow \Delta_{i,j} \cup \perp$ : Takes a secret key  $\text{sk}_i$  and public key  $\text{pk}_j$  and outputs an update token  $\Delta_{i,j}$ , or  $\perp$  when  $i = j$ . This last condition is often left out of constructions for compactness.
- $\text{ReEnc}(\Delta_{i,j}, C) \rightarrow C'$ : Takes a ciphertext  $C$  under  $\text{pk}_i$  and outputs a new ciphertext  $C'$  under  $\text{pk}_j$ .

A PRE scheme is correct if, for all  $m \in \mathcal{M}$ ,  $(\text{pk}, \text{sk}) \stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda)$ , then:

$$\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) \rightarrow m$$

and if, for all  $C \in \mathcal{C}$  such that  $\text{Dec}(\text{sk}_i, C) \rightarrow m$ , then:

$$\text{Dec}(\text{sk}_j, \text{ReEnc}(\Delta_{i,j}, C)) \rightarrow m$$

where  $(\text{pk}_i, \text{sk}_i), (\text{pk}_j, \text{sk}_j) \stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda)$  and  $\Delta_{i,j} \leftarrow \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$ .

Note that some PRE constructions have a *correctness bound* – a limit on the number of re-encryptions that are possible before the resulting ciphertext fails to decrypt properly. We shall see this in Section 6.

### 3.1 Re-encryption graphs

We often use a directed re-encryption graph (DRG) when discussing the security of PRE schemes. A DRG tracks queries the adversary  $\mathcal{A}$  makes during a security game to represent re-encryptions that  $\mathcal{A}$  can make locally. The DRG consists of nodes  $v_i$  that represent key pairs, and directed edges  $\vec{e}_{i,j}$  which are added when a  $\text{O}_{\text{ReKeyGen}}(i, j)$  query is made. Using update tokens, the adversary can locally re-encrypt challenge ciphertexts. Therefore, if a challenge ciphertext is an encryption under  $\text{pk}_i$ , and there exists a sequence of tokens going from  $i$  to  $j$ , then both  $\text{sk}_i$  and  $\text{sk}_j$  are considered challenge keys. Represented using a graph  $\text{DRG}$ , if  $\text{DRG}$  contains a path from  $v_i$  to  $v_j$  and  $\text{sk}_i$  is a challenge key, then so is  $\text{sk}_j$ . The DRG is often used to enforce the *trivial win condition* which states that the adversary cannot query oracles to reveal a challenge under a corrupted key. This is a standard condition in all PRE security definitions. Figure 1 gives a pictorial representation of this.

For example, for simply rotating keys the resulting graph will be a chain, as is assumed in *updatable encryption* [7, 22] and *key rotation for authenticated*

---

Note that some definitions of a PRE scheme have an additional input  $\ell$  to indicate a *level* the ciphertext should be at. In this work, we leave out  $\ell$  unless discussing schemes and results that use levelling explicitly.

if a scheme is bidirectional, then edges added would be directionless. In this work we mainly focus on unidirectional schemes



Figure 1: An example directed re-encryption graph,  $\mathcal{DRG}$ . If a **challenge** is learned under  $\text{pk}_6$  and  $\mathcal{A}$  has learned tokens  $\Delta_{6,5}$  and  $\Delta_{6,7}$ , then  $\text{pk}_5, \text{pk}_6$  are also considered **challenge keys**. Therefore, token queries that lead to paths from **challenge keys** to **corrupted** ones such as  $\Delta_{6,1}$  or  $\Delta_{5,2}$  result in a trivial win. The graph on the right gives some examples of what tokens the adversary can and cannot learn next.

encryption [16], whereas some access control hierarchies may lead to trees. Some results in existing work [17] mainly apply to certain types of graph. Re-encryption graphs are often used in security proofs, subject to graphs being *acyclic*. We make this assumption throughout this paper.

### 3.2 Common oracles

Most definitions use the same oracles, which we define in Figure 2 for compactness. The main variations between definitions are how the challenge oracle  $\mathcal{O}_{\text{challenge}}$  is defined, and sometimes whether  $\mathcal{O}_{\text{ReEnc}}$  affects the DRG. We therefore define these in each individual game. Games often keep track of lists updated by the oracles, namely a list of challenge keys  $\mathcal{K}_{\text{chal}}$ , corrupted keys  $\mathcal{K}_{\text{corrupted}}$ , oracle-generated ciphertexts  $\mathcal{C}_{\text{honest}}$ , challenge ciphertexts  $\mathcal{C}_{\text{chal}}$  and oracle-generated tokens  $\mathcal{T}_{\text{honest}}$ .

$\mathcal{O}_{\text{KeyGen}}(1^\lambda)$	$\mathcal{O}_{\text{Corrupt}}(i)$	$\mathcal{O}_{\text{Enc}}(i, m)$	$\mathcal{O}_{\text{ReKeyGen}}(i, j)$
$\kappa = \kappa + 1$ $(\text{pk}_\kappa, \text{sk}_\kappa) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ <div style="border: 1px solid black; padding: 2px; display: inline-block;"><math>\mathcal{DRG.add}(v_\kappa)</math></div> <b>return</b> $\text{pk}_\kappa$	$\mathcal{K}_{\text{corrupted.add}}(\text{sk}_i)$ <b>return</b> $\text{sk}_i$	$C \xleftarrow{\$} \text{Enc}(\text{pk}_i, m)$ $\mathcal{C}_{\text{honest.add}}(i, C)$ <div style="border: 1px solid black; padding: 2px; display: inline-block;"><math>\mathcal{C}_{\text{msg}}[i, C] = m</math></div> <b>return</b> $C$	<b>if</b> $\mathcal{DRG} \cup \vec{e}_{i,j}$ is cyclic: <b>return</b> $\perp$ $\Delta_{i,j} \xleftarrow{\$} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$ $\mathcal{T}_{\text{honest.add}}(i, j, \Delta_{i,j})$ <div style="border: 1px solid black; padding: 2px; display: inline-block;"><math>\mathcal{DRG.add}(\vec{e}_{i,j})</math></div> <b>return</b> $\Delta_{i,j}$

Figure 2: Common oracles used in security games for PRE.  $\kappa$  is the number of keys in the game. Boxed values indicate steps to update lists that may not be used depending on the game. The lists a particular game uses are indicated in the game's setup phase.

$\mathcal{O}_{\text{KeyGen}}(1^\lambda)$  increments the number of keys  $\kappa$ , generates a new key pair and adds a new vertex to  $\mathcal{DRG}$  to represent this key pair before returning the public key.  $\mathcal{O}_{\text{Corrupt}}(i)$  adds  $\text{sk}_i$  to the set of corrupted keys  $\mathcal{K}_{\text{corrupted}}$  and returns it



to the adversary.  $\mathcal{O}_{\text{Enc}}(i, m)$  encrypts the message under  $\text{pk}_i$ , appends the list of honestly-created ciphertexts  $\mathcal{C}_{\text{honest}}$  with  $(i, C)$  and returns the ciphertext. Sometimes we also need to consider a lookup table which can give the message used to create an honestly-generated ciphertext, to use proof techniques where challenge ciphertexts are replaced with fresh encryptions when the key is unknown. In this case,  $\mathcal{C}_{\text{msg}}$  is used.  $\mathcal{O}_{\text{ReKeyGen}}(i, j)$  creates the update token, appends the list of honestly-generated update tokens  $\mathcal{T}_{\text{honest}}$  with  $(i, j, \Delta_{i,j})$ , adds the edge  $\vec{e}_{i,j}$  to  $\mathcal{DRG}$  and returns the update token.

In our syntax, the restrictions on what tokens the adversary is allowed to learn is not enforced by oracles (as in other work), but instead by the list of challenge keys  $\mathcal{K}_{\text{chal}}$  being updated using  $\mathcal{DRG}$  at the end of the game. The following function is used to update the set of challenge keys:

```

UpdateChallengeKeys( $\mathcal{K}_{\text{chal}}, \mathcal{DRG}$ )
-----
 $\forall i$  such that  $\text{sk}_i \in \mathcal{K}_{\text{chal}}$  :
   $\forall j$  such that  $\exists$  a path from  $v_i$  to  $v_j$  in  $\mathcal{DRG}$  :
     $\mathcal{K}_{\text{chal}} \cdot \text{add } \text{sk}_j$ 
return  $\mathcal{K}_{\text{chal}}$ 

```

We enforce the trivial win condition by calling `UpdateChallengeKeys` at the end of the game, and checking that no challenge keys have been corrupted ( $\mathcal{K}_{\text{chal}} \cap \mathcal{K}_{\text{corrupted}} = \emptyset$ ).

### 3.3 Indistinguishability definitions

Here we give a number of indistinguishability notions, covering both traditional PKE and message indistinguishability in PRE. For the present we only concern ourselves with security for selective key corruptions. Security with adaptive key corruptions discussed in Section 7.

**IND-CPA in PKE and PRE.** IND-CPA-security is a well-known notion in public-key encryption which states that given a ciphertext, an adversary cannot distinguish which of two messages it is an encryption of.

$\text{PKE-CPA}_{\mathcal{A}}^{b, \mathcal{PK}\mathcal{E}}(1^\lambda)$	$\mathcal{O}_{\text{challenge}}^{\text{PKE-CPA}}(i, m_0, m_1)$
$\kappa = 0$	<b>if</b> $ m_0  \neq  m_1 $ : <b>return</b> $\perp$
$b' \leftarrow \mathcal{A}_1^{\mathcal{O}_{\text{KeyGen}}, \mathcal{O}_{\text{challenge}}^{\text{PKE-CPA}}}(1^\lambda)$	$C \xleftarrow{\$} \text{Enc}(\text{pk}_i, m_b)$
<b>return</b> $b'$	<b>return</b> $C$

Figure 3: The PKE-CPA game. This is usually called the IND-CPA game, but we make the distinction here for indistinguishability of chosen-plaintext attacks for PKE schemes and PRE schemes.

**Definition 3.** A PKE scheme  $\mathcal{PK}\mathcal{E}$  is  $\epsilon$ -Indistinguishable against Chosen Plaintext Attacks ( $\epsilon$ -PKE-CPA-secure) if for all Probabilistic Polynomial-Time (PPT) adversaries  $\mathcal{A}$ :

$$\left| \Pr \left[ \text{PKE-CPA}_{\mathcal{A}}^{0, \mathcal{PK}\mathcal{E}}(1^\lambda) = 1 \right] - \Pr \left[ \text{PKE-CPA}_{\mathcal{A}}^{1, \mathcal{PK}\mathcal{E}}(1^\lambda) = 1 \right] \right| \leq \epsilon$$

where PKE-CPA is defined in Figure 3. If  $\epsilon$  is negligible as parameterised by the security parameter  $\lambda$ , then we say the scheme is Indistinguishable against Chosen Plaintext Attacks (PKE-CPA-secure).

A PRE scheme  $\mathcal{PRE}$  is  $\epsilon$ -IND-CPA secure if the PKE scheme given by  $\mathcal{PK}\mathcal{E} = \{\mathcal{PRE}.\text{KeyGen}, \mathcal{PRE}.\text{Enc}, \mathcal{PRE}.\text{Dec}\}$  is  $\epsilon$ -PKE-CPA-secure.

**Extending IND-CPA for PRE.** The most common definition of security for PRE builds on the previous definition, stating that the scheme should still be IND-CPA secure when given the additional functionality of re-encryption. This reflects that the proxy should not learn the message during the re-encryption process. We now give a definition for CPA security for PRE schemes.

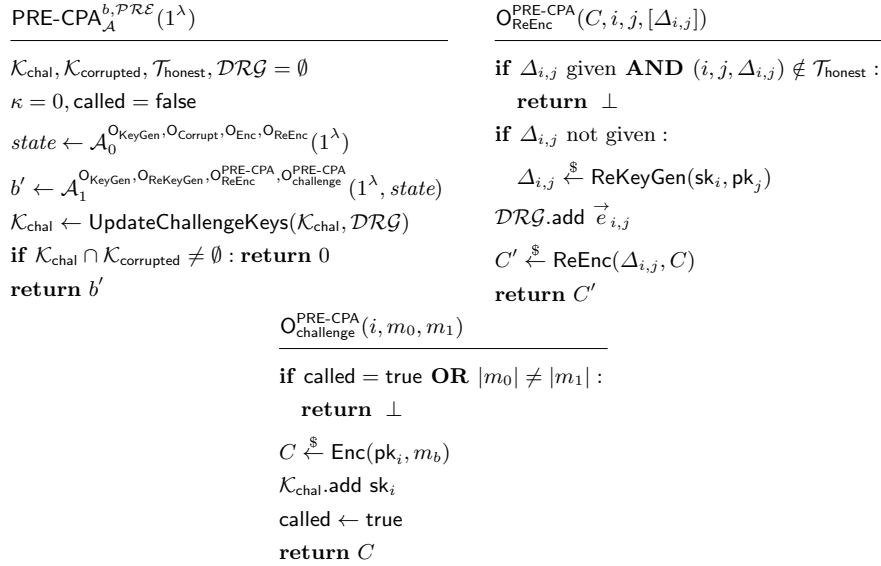


Figure 4: The PRE-CPA game – an extension of IND-CPA which accounts for re-encryption.  $\text{OKeyGen}, \text{OEnc}, \text{OReKeyGen}$  are as defined in Figure 2.

**Definition 4.** A PRE scheme  $\mathcal{PRE}$  is said to be (selectively)  $\epsilon$ -PRE-Indistinguishable against Chosen Plaintext Attacks-secure ( $\epsilon$ -PRE-CPA-secure) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$\left| \Pr \left[ \text{PRE-CPA}_{\mathcal{A}}^0(1^\lambda) = 1 \right] - \Pr \left[ \text{PRE-CPA}_{\mathcal{A}}^1(1^\lambda) = 1 \right] \right| \leq \epsilon,$$

where PRE-CPA is defined in Figure 4.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme is (selectively) PRE-Indistinguishable against Chosen Plaintext Attacks-secure (PRE-CPA-secure).

Whilst the above definition is based on the one given in [17], our formulation is slightly different as we account for there being multiple possible tokens per key pair, meaning  $\text{O}_{\text{ReEnc}}$  allows  $\mathcal{A}$  to input an honestly-generated update token as opposed to only having indexes as input. Note that the  $\mathcal{DRG}$  is created by adding an edge whenever  $\text{O}_{\text{ReEnc}}$  is called.

**Honest Re-encryption Attacks.** Recently, a stronger notion than PRE-CPA security has been introduced which allows the adversary to re-encrypt non-challenge ciphertexts to any key, as long as the ciphertexts were honestly generated. Cohen formalised these as *Honest Re-encryption Attacks (HRA)* [12] but the same idea is also used elsewhere [22]. This motivates IND-HRA-security [12]. We base our formulation on IND-ENC-security [22], HRA-security [12] and IND-CPA-security [30].

$\text{IND-HRA}_{\mathcal{A}}^{b, \mathcal{PRE}}(1^\lambda)$	
$\mathcal{K}_{\text{chal}}, \mathcal{K}_{\text{corrupted}}, \mathcal{C}_{\text{honest}}, \mathcal{C}_{\text{chal}}, \mathcal{T}_{\text{honest}}, \mathcal{DRG} = \emptyset$	
$\kappa = 0, \text{called} = \text{false}$	
$state \leftarrow \mathcal{A}_0^{\text{O}_{\text{KeyGen}}, \text{O}_{\text{Corrupt}}, \text{O}_{\text{Enc}}, \text{O}_{\text{ReEnc}}}(1^\lambda)$	
$b' \leftarrow \mathcal{A}_1^{\text{O}_{\text{KeyGen}}, \text{O}_{\text{Enc}}, \text{O}_{\text{ReKeyGen}}, \text{O}_{\text{ReEnc}}^{\text{IND-HRA}}, \text{O}_{\text{challenge}}^{\text{IND-HRA}}}(1^\lambda, state)$	
$\mathcal{K}_{\text{chal}} \leftarrow \text{UpdateChallengeKeys}(\mathcal{K}_{\text{chal}}, \mathcal{DRG})$	
<b>if</b> $\mathcal{K}_{\text{chal}} \cap \mathcal{K}_{\text{corrupted}} \neq \emptyset$ : <b>return</b> 0	
<b>return</b> $b'$	
$\text{O}_{\text{ReEnc}}^{\text{IND-HRA}}(C, i, j, [\Delta_{i,j}])$	$\text{O}_{\text{challenge}}^{\text{IND-HRA}}(i, m_0, m_1)$
<b>if</b> $(i, C) \notin \mathcal{C}_{\text{honest}}$ :	<b>if</b> $ m_0  \neq  m_1 $ <b>OR</b> $\text{called} = \text{true}$ :
<b>return</b> $\perp$	<b>return</b> $\perp$
<b>if</b> $\Delta_{i,j}$ given <b>AND</b> $(i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}}$ :	$C \stackrel{\$}{\leftarrow} \text{ReEnc}(\text{pk}_i, m_b)$
<b>return</b> $\perp$	$\mathcal{C}_{\text{honest}}.\text{add}(i, C)$
<b>if</b> $\Delta_{i,j}$ not given :	$\mathcal{C}_{\text{chal}}.\text{add}(i, C)$ $\mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_i)$
$\Delta_{i,j} \stackrel{\$}{\leftarrow} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$	$\text{called} \leftarrow \text{true}$
$C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C)$	<b>return</b> $C$
$\mathcal{C}_{\text{honest}}.\text{add}(j, C')$	
<b>if</b> $(i, C) \in \mathcal{C}_{\text{chal}}$ :	
$\mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_j)$	
<b>return</b> $C'$	

Figure 5: The IND-HRA game. Like the HRA model [12], it allows re-encryptions of non-challenge ciphertexts to compromised keys using  $\text{O}_{\text{ReEnc}}$ .

**Definition 5.** A PRE scheme  $\mathcal{PRE}$  is said to be (selectively)  $\epsilon$ -Indistinguishable against Honest Re-encryption Attacks-secure ( $\epsilon$ -IND-HRA-secure) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$|\Pr[\text{IND-HRA}_{\mathcal{A}}^0(1^\lambda) = 1] - \Pr[\text{IND-HRA}_{\mathcal{A}}^1(1^\lambda) = 1]| \leq \epsilon,$$

where  $\text{IND-HRA}_{\mathcal{A}}^{b, \mathcal{PRE}}$  is defined in Figure 5.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme is (selectively) Indistinguishable against Honest Re-encryption Attacks-secure (IND-HRA-secure).

We discuss security with respect to adaptive key corruptions in Section 7.

**Theorem 1.**  $\text{IND-HRA} \implies \text{PRE-CPA} \implies \text{PKE-CPA}$ .

As each game builds directly on the last but giving the adversary access to more information, the proof of this theorem follows trivially.

## 4 Existing Definitions that Re-randomise

Recall from the key encapsulation example that re-randomisation is necessary for PCS. The definitions covered in Section 3.3 are not designed to examine re-randomisation, and are therefore too weak for key revocation considerations. In this section, we discuss IND-UPD security for asymmetric encryption, as IND-UPD is the closest definition for PCS for PRE. We also give a definition for source-hiding which we will show in Theorem 2 gives PCS when a scheme is also IND-HRA-secure. Whilst these definitions consider key corruption to an extent, they do not consider an adversary who can learn the old secret key and update token and therefore they do not cover full compromise of the user who generated the update token.

### 4.1 Asymmetric IND-UPD

IND-UPD [22, definition 3] is a post-compromise security notion for updatable encryption schemes, which are a variant of symmetric PRE schemes. We consider a version adapted to the public-key setting, for easier comparison to our definitions. We call this pkIND-UPD, for which we give the main points here, and a full description in Appendix B.

In the pkIND-UPD game, key updates happen sequentially. The challenge oracle outputs a re-encrypted ciphertext and the adversary must guess which ciphertext it is a re-encryption of. Challenge ciphertexts are updated whenever a new key is generated, but only given to the adversary if the oracle  $\mathcal{O}_{\text{LearnChal}}^{\text{pkIU}}$  is called. One of the winning conditions is given in  $\mathcal{O}_{\text{ReEnc}}^{\text{pkIU}}$  is that when  $\text{ReEnc}$  is deterministic, the adversary cannot have re-encrypted either of the potential challenge ciphertexts  $\tilde{C}_0, \tilde{C}_1$ . Another notable condition is that the adversary cannot learn the update token going towards the key that the challenge is given

under, which is enforced as a condition in  $\mathcal{O}_{\text{LearnTok}}^{\text{pkIU}}$ . The final constraint is related to directionality, that if the scheme is bidirectional then  $\mathcal{A}$  cannot have learned any tokens leading from corrupted keys to challenge keys. We will address these points when we come to build our own definition in Section 5.

## 4.2 Source-hiding

Fuchsbauer et al. [17] define *source-hiding* as a component for demonstrating that PRE security with selective key corruptions can imply security with adaptive key corruptions in restricted re-encryption circumstances. Informally, in a source-hiding scheme it is possible to create a fresh encryption of a message that is indistinguishable from a re-encrypted ciphertext that is an encryption of the same message. This means re-encrypted ciphertexts reveal no history as to the keys they were previously encrypted under, or similarities between components of previous ciphertexts.

We give a formal description of the game defining the source-hiding property in Figure 6. Our formulation generalises the original definition in [17] by allowing the adversary to receive  $\kappa$  keypairs rather than 1. Moreover, as before, we allow the adversary to query multiple re-key tokens between any key pairs of their choice.

$\text{SH}_{\mathcal{A}}^{b, \text{PRE}}(1^\lambda, 1^\kappa, 1^L)$	$\mathcal{O}_{\text{challenge}}^{\text{SH}}(i, j, \Delta_{i,j}^*, m^*, \ell^*)$
$\{(\text{pk}_\iota, \text{sk}_\iota) \xleftarrow{\$} \text{KeyGen}(1^\lambda)\}_{\iota \in [\kappa]}$ $\mathcal{T}_{\text{honest}} = \emptyset$ $b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{challenge}}^{\text{SH}}, \mathcal{O}_{\text{ReKeyGen}}}(1^\lambda, \{(\text{pk}_\iota, \text{sk}_\iota)\}_{\iota \in [\kappa]})$ <b>return</b> $b'$	<b>if</b> $[(i, j, \Delta_{i,j}^*) \notin \mathcal{T}_{\text{honest}} \text{ OR } \ell^* + 1 > L]$ : <b>return</b> $\perp$ $C^* \leftarrow \text{Enc}(\text{pk}_i, m^*, \ell^*)$ $C^{(0)} \xleftarrow{\$} \text{ReEnc}(\Delta^*, C^*)$ $C^{(1)} \xleftarrow{\$} \text{Enc}(\text{pk}_j, m^*, \ell^* + 1)$ <b>return</b> $(C^*, C^{(b)})$

Figure 6: Experiments for the source-hiding property. Here,  $\ell$  denotes a *level* for the ciphertext to be encrypted at – essentially the number of times  $C$  has been re-encrypted. This is important for noisy PRE schemes, but ignored for PRE schemes without levelling.  $L$  is the number of times a ciphertext can be re-encrypted without breaking the correctness conditions (the correctness bound).

**Definition 6.** A PRE scheme  $\text{PRE}$  is said to be  $\epsilon$ -source-hiding ( $\epsilon$ -SH) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$\left| \Pr \left[ \text{SH}_{\mathcal{A}}^{0, \text{PRE}}(1^\lambda, 1^\kappa, 1^L) = 1 \right] - \Pr \left[ \text{SH}_{\mathcal{A}}^{1, \text{PRE}}(1^\lambda, 1^\kappa, 1^L) = 1 \right] \right| \leq \epsilon$$

where  $\text{SH}_{\mathcal{A}}^{b, \text{PRE}}$  is defined in Figure 6.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme is source-hiding (SH).

## 5 Strong PCS for PRE

In this section, we create a definition for strong PCS for public-key PRE schemes. We begin by justifying our motivation, explaining why we believe existing definitions are insufficient for PCS.

### 5.1 Motivation for a new definition

We have two main motivations for creating a new definition. The first is that there is currently no definition that implies unidirectionality and ciphertext re-randomisation. We believe that the distinction of whether a PRE scheme is unidirectional or bidirectional is vital in the post-compromise scenario to model the corruption of used update tokens, and hence we define PCS to explicitly mean that schemes meeting the definition must be unidirectional. The second motivation is a matter of existing definitions inherently assuming which of the algorithms in the PRE scheme are probabilistic. We explain how this introduces problems when considering a post-compromise scenario. We now go into both of these motivations in more detail.

**Explicit unidirectionality.** We use the most relevant definition in existing literature to make our case. IND-UPD [22] places restrictions based on inferable information, as defined by the [22] notions of directionality:

- When  $sk_j$  cannot be derived from  $sk_i$  and  $\Delta_{i,j}$  (LT-unidirectional)
- When  $sk_j$  can be derived from  $sk_i$  and  $\Delta_{i,j}$  (LT-bidirectional).

In the LT-unidirectional case, the adversary can acquire re-encryption tokens from a corrupted key to a challenge key, but not the other way around. In the LT-bidirectional case, the adversary is additionally prevented, by definition, from learning tokens from challenge keys to corrupted keys or vice versa. This means that for bidirectional schemes, the adversary queries tokens in such a way that the resulting re-encryption graphs form disjoint sub-graphs – one containing corrupted keys and the other containing challenge keys. Proving security is therefore reduced to proving that unrelated, randomly-generated keys cannot be used to infer information about an encrypted message. We consider this too restrictive for the intuition of PCS.

Since the derivability of the new secret key is implicit information to the game, meeting the definition itself says nothing about directionality and therefore whether update token compromise is of concern. We believe that it is better to have directionality given explicitly as it allows practitioners to understand the security of a scheme more clearly.

**Assuming probabilistic algorithms.** There appear to be only two existing security definitions which explicitly consider re-randomisation [16, 22]. The [16]

---

The general understanding of unidirectionality is not so strong - the new key does not necessarily have to be derivable, but the token and old key should lead to the message being learned.

definition of a *key rotation scheme* assumes that  $\text{ReKeyGen}$  is randomised but  $\text{ReEnc}$  is deterministic. This leads to a necessary condition that the update token used to create the challenge re-encryption cannot be learned by the adversary, otherwise the adversary could use it to re-encrypt the input ciphertexts locally and compare this to the challenge to win the game. The adversary is allowed to learn other tokens going from corrupted to challenge keys using an oracle  $\mathcal{O}_{\text{ReKeyGen}}$ , but not the exact token used to perform the re-encryption. This means UP-REENC [16] does not model compromise of the update token used. For this reason, it is important that new randomness is not just introduced via the update token to prevent trivial downgrading of the challenge ciphertext if the adversary compromises the update token used.

In the [22] definition of an *updatable encryption scheme*, the opposite assumption is made – that  $\text{ReEnc}$  is randomised and  $\text{ReKeyGen}$  is deterministic. This means that for keys  $\text{sk}_i, \text{pk}_j$ , there is only one update token  $\Delta_{i,j}$ . This is reflected in their IND-UPD security game (and  $\text{pkIND-UPD}$ ) by having all tokens generated at the start of the game and later having oracles reveal tokens to the adversary. This is less fitting for schemes where  $\text{ReKeyGen}$  is randomised and there are multiple tokens per key pair. More importantly, such an assumption rules out the ways in which secret keys are masked in the update token, which is important for PCS. The BV-PRE scheme is an example of this, where knowledge of the key ‘ $\text{pk}_j$ ’ together with  $\Delta_{i,j}$  can be used to derive  $\text{sk}_i$ . Another example are ElGamal-based symmetric PRE schemes (e.g. [6, 22]) where update tokens have the form  $\Delta_{i,j} = \text{sk}_j / \text{sk}_i$ . Clearly, given the update token, compromise of the old key leads to compromise of the new key. Introducing some randomness gives a means of masking the new key. It also means that the client does not need to fully rely on the proxy to be assured of new randomness, which is more appropriate for some trust scenarios.

For constructions where randomness is added in both  $\text{ReKeyGen}$  and  $\text{ReEnc}$ , neither definition is suitable. It is therefore of interest to create a security notion for PCS which factors in the possibility that both the  $\text{ReKeyGen}$  and  $\text{ReEnc}$  algorithms are probabilistic.

## 5.2 Post-Compromise Security

In our notion of Post-Compromise Security (PCS), an adversary  $\mathcal{A}$  chooses two ciphertexts and a re-encryption token, and receives a challenge ciphertext which is a re-encryption of one of the original ciphertexts created using the specified token.  $\mathcal{A}$  attempts to distinguish which ciphertext was re-encrypted. This models the compromise of all key-related material prior to the challenge re-encryption.

Unlike some previous definitions, we allow  $\mathcal{A}$  to corrupt the secret key with which the update token was generated. As in IND-HRA security, we also allow  $\mathcal{A}$  to re-encrypt honestly(oracle)-generated non-challenge ciphertexts to corrupted keys. Challenges are obtained via the challenge oracle  $\mathcal{O}_{\text{challenge}}^{\text{PC}}$  which will only

---

Interestingly, other works such as [17] take a similar approach to the adversary learning update tokens, despite their assuming randomised tokens.

accept as input honestly-generated ciphertexts of the same length, and honestly-generated update tokens. The challenger maintains lists to enforce this:  $\mathcal{C}_{\text{honest}}$  and  $\mathcal{T}_{\text{honest}}$  respectively.

Here we present a formal definition PCS for general PRE. In the first stage, the adversary  $\mathcal{A}$  can access a key corruption oracle  $\text{O}_{\text{Corrupt}}$  and in the second stage they can access the challenge oracle  $\text{O}_{\text{challenge}}^{\text{PC}}$  and re-encryption oracle  $\text{O}_{\text{ReKeyGen}}$ . As in previous definitions,  $\mathcal{K}_{\text{corrupted}}$  and  $\mathcal{C}_{\text{chal}}$  are also maintained, which record corrupted keys and challenge ciphertexts, respectively and used to enforce the trivial win condition at the end of the game.

$\text{PostComp}_{\mathcal{A}}^{b, \text{PRE}}(1^\lambda)$ <hr style="border: 0.5px solid black;"/> $\mathcal{K}_{\text{chal}}, \mathcal{K}_{\text{corrupted}}, \mathcal{C}_{\text{honest}}, \mathcal{C}_{\text{chal}}, \mathcal{C}_{\text{msg}}, \mathcal{T}_{\text{honest}}, \text{DRG} = \emptyset$ $\kappa = 0, \text{called} = \text{false}$ $\text{state} \leftarrow \mathcal{A}_0^{\text{O}_{\text{KeyGen}}, \text{O}_{\text{Corrupt}}, \text{O}_{\text{Enc}}, \text{O}_{\text{ReEnc}}^{\text{PC}}}(1^\lambda)$ $b' \leftarrow \mathcal{A}_1^{\text{O}_{\text{KeyGen}}, \text{O}_{\text{Enc}}, \text{O}_{\text{ReKeyGen}}, \text{O}_{\text{ReEnc}}^{\text{PC}}, \text{O}_{\text{challenge}}^{\text{PC}}}(1^\lambda, \text{state})$ $\mathcal{K}_{\text{chal}} \leftarrow \text{UpdateChallengeKeys}(\mathcal{K}_{\text{chal}}, \text{DRG})$ $\text{if } \mathcal{K}_{\text{chal}} \cap \mathcal{K}_{\text{corrupted}} \neq \emptyset : \text{return } 0$ $\text{return } b'$	$\text{O}_{\text{ReEnc}}^{\text{PC}}(C, i, j, [\Delta_{i,j}])$ <hr style="border: 0.5px solid black;"/> $\text{if } \Delta_{i,j} \text{ given :}$ $\quad \text{if } (i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}} : \text{return } \perp$ $\quad \text{else : } \Delta_{i,j} \leftarrow \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$ $\text{if } (i, C) \notin \mathcal{C}_{\text{honest}} : \text{return } \perp$ $C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C)$ $\mathcal{C}_{\text{honest}}.\text{add}(j, C')$ $\mathcal{C}_{\text{msg}}[(j, C')] = \mathcal{C}_{\text{msg}}[(i, C)]$ $\text{if } (i, C) \in \mathcal{C}_{\text{chal}} :$ $\quad \mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_j)$ $\text{return } C'$
$\text{O}_{\text{challenge}}^{\text{PC}}(C_0, C_1, i, j, \Delta_{i,j})$ <hr style="border: 0.5px solid black;"/> $\text{if }  C_0  \neq  C_1  \text{ OR called} = \text{true} : \text{return } \perp$ $\text{if } (i, C_0), (i, C_1) \notin \mathcal{C}_{\text{honest}} \text{ OR } (i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}} : \text{return } \perp$ $C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C_b)$ $\mathcal{C}_{\text{msg}}[(j, C')] = \mathcal{C}_{\text{msg}}[(i, C_b)]$ $\mathcal{C}_{\text{honest}}.\text{add}(j, C'), \mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_j)$ $\text{called} \leftarrow \text{true}$ $\text{return } C'$	

Figure 7: The PostComp game. This reflects full compromise of the old secret key and update token used to perform the re-encryption.

**Definition 7.** A PRE scheme  $\text{PRE}$  is said to have (selective)  $\epsilon$ -Post-Compromise Security ( $\epsilon$ -PCS) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$|\Pr[\text{PostComp}_{\mathcal{A}}^0(1^\lambda) = 1] - \Pr[\text{PostComp}_{\mathcal{A}}^1(1^\lambda) = 1]| \leq \epsilon,$$

where  $\text{PostComp}_{\mathcal{A}}^{b, \text{PRE}}$  is defined in Figure 7.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme has (selective) Post-Compromise Security (PCS).



We give a definition of PCS for symmetric PRE schemes in Appendix A, and for adaptive key corruptions in Section 7.

### 5.3 Basic observations

In the following lemma, we show that it is necessary for PCS to have randomness incorporated into re-encryption.

**Lemma 1.** *No PRE scheme where ReEnc is deterministic has PCS.*

*Proof.* If ReEnc is deterministic then  $\mathcal{A}$  can submit  $(C_0, C_1, i, j, \Delta_{i,j})$  to  $\mathcal{O}_{\text{challenge}}$  to learn challenge  $C'$ . Then  $\mathcal{A}$  can locally compute  $C'_0 \leftarrow \text{ReEnc}(\Delta_{i,j}, C_0)$  and compare this with  $C'$  – if they match then output  $b' = 0$ , otherwise output  $b' = 1$ .

**Lemma 2.** *PCS  $\implies$  unidirectional.*

*Proof.* We show that if a scheme is bidirectional then it cannot have PCS. Bidirectionality implies that an update token  $\Delta_{i,j}$ , can be used to derive  $\Delta_{j,i}$ . The adversary  $\mathcal{A}$  proceeds as follows:  $\mathcal{A}$  first corrupts a key  $\text{sk}_i$ , then creates two ciphertexts  $C_0 \leftarrow \mathcal{O}_{\text{Enc}}(i, m_0), C_1 \leftarrow \mathcal{O}_{\text{Enc}}(i, m_1)$  and update token  $\Delta_{i,j} \leftarrow \mathcal{O}_{\text{ReKeyGen}}(i, j)$ , before submitting  $(C_0, C_1, i, j, \Delta_{i,j})$  to the challenge oracle.  $\mathcal{A}$  then computes  $\Delta_{j,i}$  from  $\Delta_{i,j}$ , and uses this to compute  $C'' \leftarrow \text{ReEnc}(\Delta_{j,i}, C')$  where  $C'$  is the received challenge re-encryption to obtain the challenge ciphertext under key  $\text{pk}_i$ , before decrypting using  $\text{sk}_i$  to win the game.

This means an easy way to disprove PCS for existing schemes is to show bidirectionality.

### 5.4 Separating examples

We now demonstrate the relationship between PCS and existing security notions and constructions by means of a number of separating examples.

**Lemma 3.** *Let  $\mathcal{PRE}$  be a PRE scheme where ReKeyGen is deterministic. If  $\mathcal{PRE}$  has PCS, then it is pkIND-UPD-secure.*

*Proof.* The PostComp adversary can easily simulate the pkIND-UPD game. In oracles, whenever there is a KeyGen, the simulator calls  $\mathcal{O}_{\text{KeyGen}}$ , and whenever there is a ReKeyGen, the simulator calls  $\mathcal{O}_{\text{ReKeyGen}}$ , and otherwise follows the procedures as described with the exception of the challenge oracle. For the challenge oracle  $\mathcal{O}_{\text{challenge}}^{\text{pkIU}}$ , the adversary replaces the challenge ciphertext with the output from  $\mathcal{O}_{\text{challenge}}^{\text{PC}}(C_0, C_1, \kappa - 1, \kappa)$ . This simulates the pkIND-UPD game perfectly, and therefore the adversary will be able to win pkIND-UPD with the same advantage as for PostComp.

**Lemma 4.** *pkIND-UPD-security  $\not\Rightarrow$  PCS.*

*Proof.* Let  $\mathcal{PRE}$  be a PRE scheme where ReEnc is deterministic, and it is pkIND-UPD-secure. By Lemma 1, this scheme is not post-compromise secure.

**Lemma 5.**  $IND\text{-}HRA\text{-}security \not\Rightarrow PCS$ .

*Proof.* Let  $\overline{\mathcal{PRE}} = (\overline{\text{KeyGen}}, \overline{\text{Enc}}, \overline{\text{Dec}}, \overline{\text{ReKeyGen}}, \overline{\text{ReEnc}})$  be a IND-CPA-secure PRE scheme, and let  $(\text{symKeyGen}, \text{symEnc}, \text{symDec})$  be an IND-CPA symmetric encryption scheme. We now define a PRE scheme using the key encapsulation method as follows:

- $\text{KeyGen}(1^\lambda) \rightarrow (\text{pk}, \text{sk}) : (\text{pk}, \text{sk}) \leftarrow \overline{\text{KeyGen}}(1^\lambda)$
- $\text{Enc}(\text{pk}, m) \rightarrow (c_0, c_1) : k \xleftarrow{\$} \text{symKeyGen}(1^\lambda), c_0 \leftarrow \overline{\text{Enc}}(\text{pk}, k), c_1 \leftarrow \text{symEnc}(k, m)$
- $\text{Dec}(\text{sk}, c) \rightarrow m' : k' \leftarrow \text{Dec}(\text{sk}, c_0), m' \leftarrow \text{symDec}(k', c_1)$
- $\text{ReKeyGen}(\text{sk}_i, \text{pk}_j) \rightarrow \Delta_{i,j} : \Delta_{i,j} \leftarrow \overline{\text{ReKeyGen}}(\text{sk}_i, \text{pk}_j)$
- $\text{ReEnc}(\Delta_{i,j}, c) \rightarrow (c'_0, c'_1) : c'_0 \leftarrow \overline{\text{ReEnc}}(\Delta_{i,j}, c_0), c'_1 = c_1$

This scheme is also IND-HRA, but is not PCS. An adversary  $\mathcal{A}$  can submit two ciphertexts to the challenge oracle, and compare the second part of the challenge re-encryption with the submitted ciphertexts to win the game.

**Lemma 6.**  $PCS \not\Rightarrow IND\text{-}HRA\text{-}security$ .

*Proof.* Let  $\overline{\mathcal{PRE}} = (\overline{\text{KeyGen}}, \overline{\text{Enc}}, \overline{\text{Dec}}, \overline{\text{ReKeyGen}}, \overline{\text{ReEnc}})$  be a PRE scheme that is IND-HRA-secure and has PCS. We use it to construct the following PRE scheme:

- $\text{KeyGen}(1^\lambda) : (\text{pk}, \text{sk}) \leftarrow \overline{\text{KeyGen}}(1^\lambda)$
- $\text{Enc}(\text{pk}, m) : C \leftarrow (m, \overline{\text{Enc}}(\text{pk}, m))$
- $\text{Dec}(\text{sk}, C) : m' \leftarrow \overline{\text{Dec}}(\text{sk}, C_1)$
- $\text{ReKeyGen}(\text{sk}_i, \text{pk}_j) : \Delta_{i,j} \leftarrow \overline{\text{ReKeyGen}}(\text{sk}_i, \text{pk}_j)$
- $\text{ReEnc}(\Delta_{i,j}, C) : C'_0 \leftarrow \overline{\text{Enc}}(\text{pk}_j, 0), C'_1 \leftarrow \overline{\text{ReEnc}}(\Delta_{i,j}, C_1)$

Clearly this scheme is not IND-HRA secure, as fresh ciphertexts contain the plaintext. However the scheme has PCS, as re-encryptions  $C'_1$  will be unrelated to  $C_1$ , since  $\overline{\mathcal{PRE}}$  has PCS, and the creation of  $C'_0$  is independent of both  $C_0$  and  $\Delta_{i,j}$ .

Since PCS does not imply any security notion concerning confidentiality of the message, confidentiality definitions must be proven separately for in order to demonstrate that a PRE scheme is useful in practice.

## 5.5 PCS via source-hiding and IND-HRA

In this section we show that a PRE scheme that is both source-hiding and IND-HRA-secure also has PCS.

**Theorem 2 (main).** *Let  $\mathcal{PRE}$  be a PRE scheme that satisfies  $\epsilon_1$ -IND-HRA-security and is  $\epsilon_2$ -SH. Let  $\mathcal{A}$ ,  $\mathcal{B}$  and  $\mathcal{C}$  be PPT algorithms that are attempting to succeed in the  $\text{PostComp}_{\mathcal{A}}^{\mathcal{PRE}}$ ,  $\text{SH}_{\mathcal{B}}^{\mathcal{PRE}}$  and  $\text{IND-HRA}_{\mathcal{C}}^{\mathcal{PRE}}$  security games, respectively. Let  $\mathcal{A}$  have advantage  $\epsilon$  in  $\text{PostComp}_{\mathcal{A}}^{\mathcal{PRE}}$ . Then:*

$$\epsilon \leq 2\epsilon_2 + \epsilon_1 < \text{negl}(\lambda),$$

for a negligible function  $\text{negl}(\lambda)$ , and thus  $\mathcal{PRE}$  has PCS.

*Proof.* Let  $\text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E},b}$  refer to the experiment in Figure 7, where the choice of  $b$  is made explicit. We prove this theorem using a sequence of hybrid steps, we start with the execution of  $\mathcal{P}\mathcal{R}\mathcal{E}$  in  $\text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E},0}$ . We show that via a sequence of reductions that this situation is computationally indistinguishable from the case where  $\mathcal{A}$  witnesses the execution in  $b = 1$ .

Firstly, we define a new oracle:

---

$\overline{\text{O}_{\text{challenge}}(C_0, C_1, i, j, \Delta_{i,j})}$

---

**if**  $|C_0| \neq |C_1|$  **OR**  $\text{called} = \text{true}$  : **return**  $\perp$   
**if**  $(i, C_0), (i, C_1) \notin \mathcal{C}_{\text{honest}}$  **OR**  $(i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}}$  : **return**  $\perp$   
 $(m_0, m_1) \leftarrow (\mathcal{C}_{\text{msg}}[(i, C_0)], \mathcal{C}_{\text{msg}}[(i, C_1)])$   
**if**  $(i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}}$  **OR**  $\text{sk}_j \in \mathcal{K}_{\text{corrupted}}$  : **return**  $\perp$   
 $C' \xleftarrow{\$} \text{Enc}(\text{pk}_j, m_b)$   
 $\mathcal{C}_{\text{msg}}[(j, C')] = m_b$   
 $\mathcal{C}_{\text{honest}}.\text{add}(j, C'), \mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_j)$   
 $\text{called} \leftarrow \text{true}$   
**return**  $C'$

Let  $\text{O}_{\text{challenge},b}$  and  $\overline{\text{O}_{\text{challenge},b}}$  be the executions of the oracles where the choice of  $b \in \{0, 1\}$  is made explicit.

- **Game<sub>0</sub>**: This is the original  $\mathcal{P}\mathcal{R}\mathcal{E}$  construction in  $\text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E},0}$ .
- **Game<sub>1</sub>**: Replace outputs from the oracle  $\text{O}_{\text{challenge},0}$  with outputs from the oracle  $\overline{\text{O}_{\text{challenge},0}}$ .
- **Game<sub>2</sub>**: Replace  $\overline{\text{O}_{\text{challenge},0}}$  with  $\overline{\text{O}_{\text{challenge},1}}$ .
- **Game<sub>3</sub>**: Replace  $\overline{\text{O}_{\text{challenge},1}}$  with  $\text{O}_{\text{challenge},1}$ .

It is not hard to see that **Game<sub>3</sub>** is identical to the execution in the case of  $\text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E},1}$ . Therefore, if we can bound the advantage in distinguishing the game transitions above by a negligible function, then the probability of distinguishing in  $\text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E},b}$  is also negligible.

**Lemma 7.** *The advantage in distinguishing **Game<sub>0</sub>** and **Game<sub>1</sub>** is bounded by  $\epsilon_2$ .*

*Proof.* This proof holds in the case where  $\mathcal{P}\mathcal{R}\mathcal{E}$  has a maximum of  $\kappa = \text{poly}(\lambda)$  nodes.  $\mathcal{B}$  receives key pairs  $(\text{pk}_\iota, \text{sk}_\iota)_{\iota \in [\kappa]}$  and re-encryption tokens  $\{\Delta_{i,j}\}_{(i,j) \in [\kappa]^2}$  for  $\mathcal{P}\mathcal{R}\mathcal{E}$ . Let  $\overline{\text{O}_{\text{challenge}}}$  be the challenge oracle that  $\mathcal{B}$  has access to in  $\text{SH}_{\mathcal{B}}^{\mathcal{P}\mathcal{R}\mathcal{E},c}$ .

Then  $\mathcal{B}$  instantiates the oracles  $\text{O}_{\text{KeyGen}}, \text{O}_{\text{Corrupt}}, \text{O}_{\text{Enc}}, \text{O}_{\text{ReKeyGen}}, \text{O}_{\text{ReEnc}}$  in the same way as in  $\text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E},b}$ . Let  $\mathcal{C}_{\text{honest}}$  and  $\mathcal{C}_{\text{msg}}$  be the sets that  $\mathcal{B}$  keeps track of the queries made to  $\text{O}_{\text{Enc}}$  and  $\text{O}_{\text{ReEnc}}$ . For a query  $(C_0^*, C_1^*, (i^*, j^*), \Delta_{i^*,j^*}^*)$  made to the oracle  $\text{O}_{\text{challenge},0}$  by  $\mathcal{A}$ , then  $\mathcal{B}$  does the following:

- **if**  $|C_0^*| \neq |C_1^*|$  : **return**  $\perp$ ;
- **if**  $(i^*, C_0^*), (i^*, C_1^*) \notin \mathcal{C}_{\text{honest}}$  : **return**  $\perp$ ;
- **if**  $(i^*, j^*, \Delta_{i^*,j^*}^*) \notin \mathcal{T}_{\text{honest}}$  : **return**  $\perp$ ;

- **if**  $\text{sk}_j \in \mathcal{K}_{\text{corrupted}}$  : **return**  $\perp$ ;
- $(m_0^*, m_1^*) \leftarrow \mathcal{C}_{\text{msg}}$ ;
- $C_0^\dagger \xleftarrow{\$} \overline{\text{O}_{\text{challenge}}}(1^\lambda, (i^*, j^*), (m_0^*, C_0^*))$ ;
- $\mathcal{C}_{\text{honest}}.\text{add}(j, C_0^\dagger)$ ;  $\mathcal{C}_{\text{msg}}[(j, C_0^\dagger)] = m_0^*$ ;  $\mathcal{C}_{\text{chal}}.\text{add}(j, C_0^\dagger)$ ;
- **return**  $C_0^\dagger$ .

In the case of  $c = 0$ , then  $C_0^\dagger \xleftarrow{\$} \text{ReEnc}(\Delta_{i,j}, C_0^*)$ . In the case of  $c = 1$ , then  $\overline{\text{O}_{\text{challenge}}}$  returns  $C_0^\dagger \xleftarrow{\$} \text{Enc}(\text{pk}_{j^*}, m_0^*)$ . Consequently, when  $c = 0$  then  $\mathcal{B}$  simulates  $\text{O}_{\text{challenge},0}$ ; for  $c = 1$  then  $\mathcal{B}$  simulates  $\overline{\text{O}_{\text{challenge},0}}$ . Now, let  $\Pr\left[0 \xleftarrow{\$} \mathcal{A}[\text{Game}_d]\right] = \epsilon_d$  be the probability that  $\mathcal{A}$  outputs  $b' = 0$  for  $d \in \{0, 1\}$ . Then:

$$|\Pr\left[0 \xleftarrow{\$} \mathcal{A}[\text{Game}_0]\right] - \Pr\left[0 \xleftarrow{\$} \mathcal{A}[\text{Game}_1]\right]| = |\delta_0 - \delta_1| = \delta$$

represents the differences in the probabilities in the two games. Then, let  $\epsilon_0$  be the probability that  $\mathcal{A}$  outputs 0 in the simulation of  $\mathcal{B}$ . If  $\epsilon_0 \approx \delta_0$  then  $\mathcal{B}$  outputs  $c' = 0$  to  $\text{SH}_{\mathcal{B}}^{\text{PRE},c}$ . Otherwise, if  $\epsilon_0 \approx \delta_1$  then  $\mathcal{B}$  outputs  $c' = 1$ .

Notice that if  $\delta$  is noticeably large, then  $\mathcal{B}$  can win the  $\text{SH}_{\mathcal{B}}^{\text{PRE},c}$  with the same advantage. In other words,  $\mathcal{B}$  can run sequential subroutines using the adversary  $\mathcal{A}$  and build up a probability distribution based on the answers of  $\mathcal{A}$ . If the answers correlate more to the probability inferred by  $\delta_0$  then  $\mathcal{B}$  answers 0, and otherwise 1.

Therefore,  $\delta < \text{Adv}^{\text{SH}}$ ; where  $\text{Adv}^{\text{SH}}$  is the advantage of  $\mathcal{B}$  in SH. By the assumption that  $\text{PRE}$  is  $\epsilon_2$ -SH, then we have that  $\delta < \epsilon_2$ .

**Lemma 8.** *The advantage in distinguishing  $\text{Game}_1$  and  $\text{Game}_2$  is bounded by  $\epsilon_1$ .*

*Proof.* Let  $\overline{\text{O}_{\text{challenge}}}$  be the challenge oracle  $\text{O}_{\text{challenge}}$  that  $\mathcal{B}$  has access to in  $\text{IND-HRA}_{\mathcal{B}}^{\text{PRE},b}(1^\lambda)$ . Firstly,  $\mathcal{B}$  simulates the oracles  $\text{O}_{\text{KeyGen}}$ ,  $\text{O}_{\text{Corrupt}}$ ,  $\text{O}_{\text{Enc}}$ ,  $\text{O}_{\text{ReKeyGen}}$ ,  $\text{O}_{\text{ReEnc}}$  using the oracles in  $\text{IND-HRA}_{\mathcal{B}}^{\text{PRE},c}(1^\lambda)$ . Then, for a query  $(C_0^*, C_1^*, (i^*, j^*), \Delta_{i,j}^*)$  made to  $\overline{\text{O}_{\text{challenge},b}}$  by  $\mathcal{A}$ ,  $\mathcal{B}$  does the following:

- **if**  $|C_0^*| \neq |C_1^*|$  : **return**  $\perp$ ;
- **if**  $(i^*, C_0^*), (i^*, C_1^*) \notin \mathcal{C}_{\text{honest}}$  : **return**  $\perp$ ;
- **if**  $(i^*, j^*, \Delta_{i^*,j^*}^*) \notin \mathcal{T}_{\text{honest}}$  : **return**  $\perp$ ;
- **if**  $\text{sk}_j \in \mathcal{K}_{\text{corrupted}}$  : **return**  $\perp$ ;
- $(m_0^*, m_1^*) \leftarrow \mathcal{C}_{\text{msg}}$ ;
- $C^\dagger \xleftarrow{\$} \overline{\text{O}_{\text{challenge}}}(m_0^*, m_1^*, j^*)$ ;
- $\mathcal{C}_{\text{msg}}[(j, C^\dagger)] = m_b^*$ ;
- **return**  $C^\dagger$ .

---

Note that  $\mathcal{A}$  has no access to  $\mathcal{C}_{\text{msg}}$  and thus the use of this set has no impact on the simulation of  $\text{O}_{\text{challenge},0}$ .

When  $b = 0$ , then  $C^\dagger \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}_j, m_0^*)$ ; otherwise, if  $b = 1$ , then  $C^\dagger \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}_j, m_1^*)$ . It is now simply the case that  $\mathcal{B}$  simulates  $\overline{\text{O}}_{\text{challenge},b}$  depending on the value of  $b$ . Thus,  $\mathcal{B}$  simply outputs  $b' = b_{\mathcal{A}}$  where  $b_{\mathcal{A}} \leftarrow \mathcal{A}$  is output by  $\mathcal{A}$ . As a result, let  $\delta$  be the advantage of  $\mathcal{A}$  in distinguishing the scenarios in both games. Then  $\delta$  is bounded by the advantage that  $\mathcal{B}$  has in  $\text{IND-HRA}_{\mathcal{B}}^{\mathcal{PRE},b}$ . By the assumption that  $\mathcal{PRE}$  satisfies  $\epsilon_1$ -IND-HRA security, then  $\delta < \epsilon_1$  and the lemma follows.

**Lemma 9.** *The advantage in distinguishing  $\text{Game}_2$  and  $\text{Game}_3$  is bounded by  $\epsilon_2$ .*

*Proof.* The simulation for  $\mathcal{A}$  by  $\mathcal{B}$  is the same as in Lemma 7 except in reverse. That is,  $\text{Game}_2$  corresponds to the case where  $c = 1$  and  $\text{Game}_3$  to the case where  $c = 0$ . The rest of the details are covered by the previous simulation and so we refer the reader to this proof for the full details.

**Lemma 10.** *The advantage of  $\mathcal{A}$  in distinguishing the executions of  $\text{PostComp}_{\mathcal{A}}^{\mathcal{PRE},b}$  in  $\text{Game}_3$  is 0 for  $b \stackrel{\$}{\leftarrow} \{0, 1\}$ .*

*Proof.* In  $\text{Game}_3$ , the sampling of  $b$  is independent of the values that are returned. In particular, the ciphertext  $C^\dagger$  returned by  $\text{O}_{\text{challenge},b}$  is always a re-encryption of the right input ciphertext  $C_1^*$ . Therefore, there is no plausible advantage that  $\mathcal{A}$  has to gain in distinguishing the two games and the lemma follows.

To complete the proof of the theorem, we note that the advantage,  $\epsilon$ , that  $\mathcal{A}$  has in distinguishing  $\text{PostComp}_{\mathcal{A}}^{\mathcal{PRE},b}$  for  $b \stackrel{\$}{\leftarrow} \{0, 1\}$  is bounded by  $2\epsilon_2 + \epsilon_1$ . Providing that there is negligible function s.t.  $\epsilon_1, \epsilon_2 < \text{negl}(\lambda)$ , then we have that  $\epsilon < \text{negl}(\lambda)$  and the proof of Theorem 2 is complete.

**Theorem 3 (main).** *Let  $\mathcal{PRE}$  be a PRE scheme which is both PKE-CPA-secure and source-hiding. Then  $\mathcal{PRE}$  also has PCS.*

*Proof.* It has been shown that PKE-CPA-security and source-hiding imply IND-HRA-security [17, Theorem 6]. This together with Theorem 2 gives us the result. For results on the precise loss of security, see [17].

**Existing PRE schemes that satisfy PCS.** Using the result of Theorem 2, we can show that a number of existing PRE schemes satisfy PCS immediately. This is advantageous since it shows that PCS is not a vacuous security model in the sense that it is achievable via well-known techniques.

Specifically, any PRE scheme that satisfies PRE-CPA-security along with source-hiding is immediately a scheme that satisfies PCS. Therefore, the schemes of [17, Construction 2, Construction 4, Construction 7.b] all satisfy PCS. Constructions 2 and 4 of [17] are based on assumptions over groups with bilinear maps; whereas Construction 7.b is based on the hardness of the decision learning with errors problem with sub-exponential noise-to-modulus ratio.

## 6 An Efficient Construction from Lattices

We introduce a natural construction with PCS, based on BV-PRE – the ring-LWE (RLWE) construction given in [30]. Whilst Theorem 3 shows that source-hiding can lead to PCS, the existing constructions with this property [17] make making sub-optimal parameter choices that significantly impact the schemes’ practicality. Our construction has PCS but is not source-hiding, implying that source-hiding is not necessary for PCS. This means that our construction can make much better parameter choices in terms of efficiency. Our construction also doesn’t rely on strong assumptions or heavy techniques such as obfuscation [11, 19]. We also achieve *transparency*, which means decryption is the same regardless of how many times the ciphertext has been re-encrypted, and the cost of decryption does not grow for repeatedly re-encrypted ciphertexts. If extra computation is needed to decrypt a re-encrypted ciphertext, then this may outweigh the benefits of outsourcing re-encryption. This fits better with motivations for outsourcing re-encryption, as heavier computation may go against the reasons for outsourcing re-encryption to begin with.

Our construction makes some adaptations to BV-PRE to fit the workflow of PRE; making use of the key resampling technique of [8] to re-randomise the ciphertext. Any scheme that permits similar re-randomisation can be proven secure using related methods. We begin this section by covering some necessary preliminaries for lattices.

### 6.1 Lattice preliminaries

We let  $q \in \mathbb{Z}$  denote some modulus and  $n$  denote a ring dimension. Furthermore, we represent the set of integers modulo  $q$  as  $\mathbb{Z}_q = \{-q/2, \dots, 0, \dots, q/2\}$ . We will be working over power-of-two cyclotomic rings of the form  $R_q = \mathbb{Z}_q[x]/(x^n + 1)$ . Next we discuss the various distributions that will be used. We use the notation  $s \stackrel{\$}{\leftarrow} D$  to denote that the element  $s$  is sampled according to distribution  $D$ . If  $D$  is a set, then we assume  $s \stackrel{\$}{\leftarrow} D$  means that  $s$  is sampled uniformly from the set  $D$ . We denote the discrete Gaussian distribution over  $\mathbb{Z}_q$  as  $\chi_\sigma$ . The distribution  $\chi_\sigma$  has its support restricted to  $\mathbb{Z}_q$  and a probability mass function proportional to that of a Gaussian distribution with variance  $\sigma^2$ . We sometimes write  $\chi_e$  where we use a subscript  $e$  to denote an “error” distribution, but the underlying variance is still denoted by  $\sigma$ . Slightly abusing notation, we can sample a polynomial  $s \stackrel{\$}{\leftarrow} \chi_e$  by sampling each of the coefficients of  $s$  according to the distribution  $\chi_e$ . We say a distribution  $D$  is  $(B, \delta)$ -bounded if  $\Pr(|x| > B : x \stackrel{\$}{\leftarrow} D) \leq \delta$ .

*Ring-RLWE (RLWE) assumption:* Let  $s$  be some secret polynomial in  $R_q$ . Samples from the  $\text{RLWE}_{n,q,\chi_e}(s)$  distribution take the form  $(a, b = as + e) \in R_q \times R_q$  where  $a \stackrel{\$}{\leftarrow} R_q, e \stackrel{\$}{\leftarrow} \chi_e$ . Note that  $\chi_e$  is referred to as the error distribution. The (normal form)  $\text{RLWE}_{n,q,\chi_e}$  problem is to distinguish between an oracle that outputs samples from  $\text{RLWE}_{n,q,\chi_e}(s)$  where  $s \stackrel{\$}{\leftarrow} \chi_e$  and an oracle that outputs

uniform elements in  $R_q \times R_q$ . The  $\text{RLWE}_{n,q,\chi_e}$  assumption states that no PPT algorithm can solve the  $\text{RLWE}_{n,q,\chi_e}$  problem with a non-negligible advantage. Note that if we take  $\sigma \geq \omega(\log n)$  and  $\sigma/q = 1/\text{poly}(n)$ , the  $\text{RLWE}_{n,q,\chi_e}$  problem is at least as hard as solving standard worst-case lattice problems over ideal lattices up to polynomial approximation factors using quantum algorithms [25].

## 6.2 Adapting BV-PRE for PCS

The underlying scheme, BV-PRE [30], is based on the BV-encryption scheme [8], which is based on RLWE. This scheme is parameterised by ciphertext modulus  $q$ , plaintext modulus  $p \geq 2$ , ring dimension  $n$ , polynomial ring  $R_q = \mathbb{Z}_q[n]/\langle x^n + 1 \rangle$  and relinearisation window  $r$ . BV-PRE is not fully public-key, relying on an additional ‘public’ key  $\text{pk}'_B$  for the target key  $s_B$  to generate update tokens. However, this key together with the token can be used to derive the old secret key. We get around this problem using the key resampling technique  $\text{ReSample}$  [8] which takes a public key  $\text{pk}_B$  and outputs a fresh public key  $\text{pk}'_B$  with the same underlying secret. We also use same relinearisation technique as [30] to reduce error growth. We refer readers interested in the full BV-PRE and re-sampling construction to Appendix C.

KeyGen( $1^\lambda$ )	Enc(pk, m)	Dec(sk, c)
$a \xleftarrow{\$} \mathcal{U}_q,$	$(a, b) \leftarrow \text{pk}$	$s \leftarrow \text{sk}$
$s, e \xleftarrow{\$} \chi_e$	$v, e_0, e_1 \xleftarrow{\$} \chi_e$	$m' = c_0 - s \cdot c_1 \pmod p$
$b = a \cdot s + pe$	$c_0 = b \cdot v + pe_0 + m$	<b>return</b> $m'$
$\text{sk} = s$	$c_1 = a \cdot v + pe_1$	
$\text{pk} = (a, b)$	<b>return</b> $c = (c_0, c_1)$	
<b>return</b> (pk, sk)		
ReKeyGen(sk <sub>A</sub> , pk <sub>B</sub> )	ReEnc( $\Delta_{A \rightarrow B}$ , pk <sub>B</sub> , c)	
<b>for</b> $i \in \{0, 1, \dots, \lfloor \log_2(q)/r \rfloor\}$ :	$\{(\beta_i, \gamma_i)\}_{i=0}^{\lfloor \log_2(q)/r \rfloor} \leftarrow \Delta_{A \rightarrow B}$	
$(\beta_i, \theta_i) \xleftarrow{\$} \text{ReSample}(\text{pk}_B)$	$(c_0, c_1) \leftarrow c$	
$\gamma_i = \theta_i - \text{sk}_A \cdot (2^r)^i$	$(\beta^{\text{proxy}}, \theta^{\text{proxy}}) \leftarrow \text{ReSample}(\text{pk}_B)$	
$\Delta_{A \rightarrow B} = \{(\beta_i, \gamma_i)\}_{i=0}^{\lfloor \log_2(q)/r \rfloor}$	$c'_0 = c_0 + \sum_{i=0}^{\lfloor \log_2(q)/r \rfloor} (c_1^{(i)} \cdot \gamma_i) + \theta^{\text{proxy}}$	
<b>return</b> $\Delta_{A \rightarrow B}$	$c'_1 = \sum_{i=0}^{\lfloor \log_2(q)/r \rfloor} (c_1^{(i)} \cdot \beta_i) + \beta^{\text{proxy}}$	
	<b>return</b> $c' = (c'_0, c'_1)$	

Figure 8: pcBV-PRE. This adapts the BV-PRE scheme [30] to be fully public-key, and minimises computation and bandwidth for the user for the re-encryption process. The key resampling algorithm  $\text{ReSample}$  is described in Appendix C.

We give our construction, pcBV-PRE, in Figure 8. It builds on BV-PRE in that randomness is also added by the proxy in the ReEnc operation. Recall that this is necessary for a scheme to have PCS, as otherwise an adversary could re-encrypt locally to obtain the same re-encryption. This change has minor implications for how this affects the correctness requirement for multiple re-encryptions over that given in [30]. Note also that pcBV-PRE inherits the IND-CPA-security proven in [30]. This is proven by following the same exact steps from the proof in [30], so we omit considerations of IND-CPA-security for brevity.

### 6.3 Security proofs

For brevity, we defer the correctness analysis of pcBV-PRE to Appendix C.

**Post-Compromise Security.** We firstly show that pcBV-PRE satisfies PCS with a direct proof. In other words, we do not leverage the proof of Theorem 2 to prove PCS via source-hiding security. This is because pcBV-PRE as written does not satisfy source-hiding security, see Section 7.3 for more details.

**Theorem 4.** *pcBV-PRE has Post-Compromise Security. In other words, for any adversary  $\mathcal{A}$  to the PostComp game,*

$$\left| \Pr \left[ \text{PostComp}_{\mathcal{A}}^{0, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda) = 1 \right] - \Pr \left[ \text{PostComp}_{\mathcal{A}}^{1, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda) = 1 \right] \right| \leq \epsilon,$$

for some  $\epsilon = \text{negl}(\lambda)$  under the  $RLWE_{n,q,\chi_e}$  assumption.

*Proof Overview.* The overall structure of the proof is essentially the same as the IND-CPA-security proof of the BV-PRE in [30]. We only give a brief informal argument here and refer the reader to the Appendix C for a full proof. The proof follows a sequence of game hops beginning with the  $\text{PostComp}^{b, \mathcal{P}\mathcal{R}\mathcal{E}}$  security game where  $b \stackrel{\$}{\leftarrow} \{0, 1\}$ . In this game, the adversary is challenged to guess the bit  $b$ . Suppose that there are  $N$  honest entities that are not corrupted by the adversary. Then we begin by making  $N$  game hops. Each game hop replaces:

1. the public key of a single honest entity with a uniform random value.
2. the re-encryption keys created using the honest entity's public key with uniform random values.

In the final game hop, the challenge ciphertext given to the adversary is a uniformly sampled value and thus the adversary has no advantage in this game. This implies that  $\text{PostComp}^{0, \mathcal{P}\mathcal{R}\mathcal{E}}$  and  $\text{PostComp}^{1, \mathcal{P}\mathcal{R}\mathcal{E}}$  are indistinguishable.

**IND-CPA security.** We secondly show that pcBV-PRE is a IND-CPA-secure PRE scheme. This is a much simpler proof since it follows a similar argument to that of Theorem 4, and the fact that the same security notion was shown in [30, Theorem 2].

**Theorem 5.** *pcBV-PRE is IND-CPA secure.*



*Proof.* The proof of this property follows a very similar argument to that of Theorem 4. Firstly, we employ the hybrid argument of Theorem 4 up to the hybrid before the final game hop is made. At this point in the security proof, the public key of all honest entities are replaced with uniform values and similarly for the re-encryption keys. The IND-CPA requirement requires that the challenge ciphertext is encrypted with respect to an honest entity. Therefore, any call to  $O_{\text{ReKeyGen}}$  is answered with a uniform response, and this is the same for responses from  $O_{\text{ReEnc}}$ .

In this case, pcBV-PRE is almost identical to the original BV-PRE scheme, since the KeyGen, Enc and Dec algorithms are defined in the exactly the same way. Moreover, the outputs of the additional algorithms are uniform using the hybrid approach above. Therefore, we can leverage the proof of [30, Theorem 2] showing that BV-PRE is IND-CPA-secure. Notice that this is possible since the oracles above can be instantiated with uniform responses, and the other oracles can be handled as in the reduction of [30]. This completes the proof.  $\square$

## 7 From selective to adaptive security

Thus far, we have discussed selective security, where the adversary first corrupts keys, then learns challenges in two distinct stages. In the adaptive model, a single-stage adversary can corrupt secret keys at any point and therefore choose which keys to corrupt as a result of received challenges. As long as the trivial win condition holds, the adversary can adaptively decide which keys should be corrupted depending on the output of other oracle queries.

To lift to this stronger model, we need notions of key privacy. The intuition behind this is that if we can replace some challenge-related queries with counterparts under different keys, and the adversary cannot tell that this has happened, then the advantage related to specific corruptions in the adaptive model is negligible.

### 7.1 Weak Key Privacy

In this section, we give the formal definitions for the related notions of key privacy. *key privacy* (sometimes called *key anonymity*), was first introduced for PKE in [4]. Informally it means that an adversary is unable to determine which public key was used to encrypt a ciphertext. A strong version for PRE schemes which accounts for update tokens appears in [1] which is referred to as *strong key privacy* in [17], and states that an adversary is unable to distinguish between an update token between two uncorrupted keys and a random element of the token space. Note that this is strictly stronger than key inference where the hidden secret keys must be computed, as is the concern in *collusion attacks*. We refer the reader to [1] for further details.

In their work relating selective and adaptive security, Fuchsbauer et al. [17] define *weak key privacy* and show that this is sufficient for adaptive security. We now define this formally.

$\text{weakKP}_{\mathcal{A}}^{b, \mathcal{PRE}}(1^\lambda, 1^\kappa)$	$\text{ReKeyGen}^{\text{wKP}}(\text{pk}_i)$
$(\text{pk}_0, \text{sk}_0), \dots, (\text{pk}_\kappa, \text{sk}_\kappa) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	$(\text{pk}'_0, \text{sk}'_0) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$
<b>for</b> $i \in \{1, \dots, \kappa\}$ :	$\Delta_{0,i}^{(1)} \leftarrow \text{ReKeyGen}(\text{sk}'_0, \text{pk}_i)$
$\Delta_{0,1}^{(0)} \leftarrow \text{ReKeyGen}(\text{sk}_0, \text{pk}_i)$	<b>return</b> $\Delta_{0,i}^{(1)}$
$\Delta_{0,1}^{(1)} \leftarrow \text{ReKeyGen}^{\text{wKP}}(\text{pk}_i)$	
$b' \leftarrow \mathcal{A}(1^\lambda, \text{pk}_0, \dots, \text{pk}_\kappa, \Delta_{0,1}^{(b)}, \dots, \Delta_{0,\kappa}^{(b)})$	
<b>return</b> $b'$	

Figure 9: The weak key privacy game. Given an update token, an adversary cannot distinguish the source public key.

**Definition 8.** A PRE scheme  $\mathcal{PRE}$  is said to have  $(\epsilon, \kappa)$ -weak key privacy if for all PPT adversaries  $\mathcal{A}$ :

$$|\Pr[\text{weakKP}_{\mathcal{A}}^0(1^\lambda, 1^\kappa) = 1] - \Pr[\text{weakKP}_{\mathcal{A}}^1(1^\lambda, 1^\kappa) = 1]| \leq \epsilon,$$

where  $\text{weakKP}_{\mathcal{A}}^{b, \mathcal{PRE}}$  is defined in Figure 9.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme has  $\kappa$ -weak key privacy. If this holds for all polynomial  $\kappa$  with respect to the security parameter, then the scheme has weak key privacy.

## 7.2 Adaptive Post-Compromise Security

In this section, discuss PCS with adaptive key corruptions. In an adaptive game, the adversary is in one stage and has access to all oracles at once, and so is free to choose which keys to corrupt based on previous oracle responses. We give the full adaptive game in Appendix E.

In general, selective security can be shown to imply adaptive security at an exponential loss, where the adversary in the selective game must first guess which keys the adaptive adversary will corrupt. This is known as *complexity leveraging*. A general framework for showing relations between selective and adaptive security for general definitions is given in [20]. Jafargholi et al. achieve this by constructing a variant of the selective game, then using pebbling games and a series of hybrids to the adaptive game to demonstrate the loss in security between each hybrid. Using pebbling games can give a smaller loss of security. A similar approach specific to PRE is used in [17], limited to when the adversary creates directed re-encryption graphs which are acyclic and have only one source node. An even tighter reduction is given which results in quasi-polynomial loss between selective and adaptive HRA security for PRE schemes with ciphertext indistinguishability and weak key privacy, limited to some types of directed re-encryption graph, namely trees and chains.

Selective PRE-CPA security combined with weak key privacy gives adaptive PRE-CPA security, where the loss of security depends on the number of keys  $\kappa$  generated during the game, and the space and time complexities of the  $\mathcal{DRG}$ .

More precisely, the loss of security is approximately  $\tau \cdot \kappa^\sigma$ , where  $\sigma$  is the maximum number of pebbles and  $\tau$  is the maximum number of moves for a valid pebbling strategy for a considered class of re-encryption graphs  $\mathcal{DRG}$  [17, Theorem 1, Theorem 5]. This means that the loss is exponential for general re-encryption graphs, but quasi-polynomial for trees and chains.

**Theorem 6 (main).** *Let  $\mathcal{PRE}$  be a PRE scheme which is  $\epsilon_1$ -selectively PRE-CPA secure,  $\epsilon_2$ -source-hiding and  $\epsilon_3$ -weakly key private. Then it is  $\epsilon$ -adaptively PCS secure with a security loss of  $\approx \tau \cdot \kappa^\sigma$  for directed re-encryption graphs  $\mathcal{DRG}$  with degree  $\kappa$ , space complexity  $\sigma$  and time complexity  $\tau$ , where*

$$\epsilon \leq 2\kappa(\kappa - 1)(Q_E + Q_{RE})Q_{RE} \cdot \epsilon_{SH} + \kappa^{\sigma+\delta+1}(\epsilon_{IND} + 2\tau \cdot \epsilon_3) \quad (1)$$

*restricted to graphs  $\mathcal{DRG}$  which are acyclic, have at most  $\kappa$  nodes, degree  $\delta$  and depth  $d$ .*

The proof of this follows from Theorem 3 and uses the same strategies as those used to prove [17, Theorem 5].

### 7.3 Achieving adaptive security of RLWE-based construction

The construction given in Figure 8 is not source-hiding. This is because fresh encryptions will have different noise magnitudes compared with re-encryptions. One method of overcoming this is using a noise ‘blurring’ approach [10], this was noted by [17, Construction 7.b]. This would involve ‘blurring’ or ‘drowning’ the noise in fresh encryptions so that it was distributed in the same way as re-encryptions. To do this, we would modify the `Enc` and `ReEnc` algorithms so that we would add fresh noise to them. This fresh noise would be taken from an error distribution such that it ‘blurred’ out all the old noise that was introduced from the encryption and re-encryption operations.

The advantage of doing this, is that on decryption the noise will not reveal anything about the method of encryption/re-encryption that was used, or the keys the ciphertext was previously encrypted under. This is vital since in the source-hiding security property, the adversary receives all of the available decryption keys and thus can decrypt any ciphertext that they want.

With this in mind, there are a number of reasons why proving source-hiding is actually a hindrance. Using the same approach as [17] requires at least a sub-exponential noise-to-modulus ratio which considerably harms performance and security. This is because the LWE assumption (and respectively RLWE in our case) that is used becomes much stronger, since the approximation factors of the related ideal lattice problems become sub-exponential rather than polynomial in  $n$ . In particular, there are quantum polynomial time algorithms solving ideal lattice problems with approximation factor  $\exp(\tilde{O}(\sqrt{n}))$  [5, 14, 15] providing evidence that we cannot simply use an arbitrary noise-to-modulus ratio while retaining the same security guarantees. Moreover, the increase in modulus that is required makes standard operations much slower. As a by-product, the scheme of [17] allows only a constant number of re-encryptions.

The intention with our construction was to give a practical PCS PRE scheme with minimal restrictions and from weak assumptions. Since our construction is very close to the original [30] definition, we achieve this goal. Since BV-PRE is comparatively fast (see [30] for exact figures) and our construction only adds extra sampling, loss of efficiency is minimal. If we wanted to incorporate source-hiding, this would result in a scheme that was impractical, based on much stronger assumptions and also heavily restricted in the number of re-encryptions that can occur. Therefore, we choose not to give an explicit formulation.

We prove that our scheme has weak key privacy in Appendix F. Therefore, it should be noted that a source-hiding version of our scheme would achieve adaptive security (for the restricted graphs of [17]) by the proof of Theorem 2 and the fact that our scheme is IND-CPA secure [5]. This result was explicitly shown in [17, Lemma 7].

## 8 Conclusions and Future work

In this paper, we have presented the strongest notion of Post-Compromise Security for PRE to date. By strongest, we mean that existing PCS notions are implied by our notion of security, and separating examples showing that the opposite implication does not follow exist. We have also shown that PCS can be achieved via a number of existing PRE security notions which immediately shows that there are existing PRE schemes that satisfy PCS [17]. Finally, we give a practical construction of a PCS secure PRE scheme with transparency which is based on lattices and discussed the possibility of achieving adaptive security for restricted re-encryption graphs. We leave as future work the possibility of proving tighter bounds between security notions, and further investigating the relationship between selective and adaptive security for more generic graphs.

## Acknowledgements

Special thanks to Katriel Cohn-Gordon for his help in motivating this work.

## References

1. Giuseppe Ateniese, Karyn Benson, and Susan Hohenberger. Key-private proxy re-encryption. In Marc Fischlin, editor, *Topics in Cryptology - CT-RSA 2009, The Cryptographers' Track at the RSA Conference 2009, San Francisco, CA, USA, April 20-24, 2009. Proceedings*, volume 5473 of *Lecture Notes in Computer Science*, pages 279–294. Springer, 2009.
2. Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Trans. Inf. Syst. Secur.*, 9(1):1–30, 2006.
3. Elaine Barker, William Barker, William Burr, William Polk, and Miles Smid. Recommendation for key management part 1: General (revision 3). *NIST special publication*, 800(57):1–147, 2012.

4. Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In Colin Boyd, editor, *Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings*, volume 2248 of *Lecture Notes in Computer Science*, pages 566–582. Springer, 2001.
5. Jean-François Biase and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 893–902. SIAM, 2016.
6. Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In Kaisa Nyberg, editor, *Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding*, volume 1403 of *Lecture Notes in Computer Science*, pages 127–144. Springer, 1998.
7. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 410–428. Springer, 2013.
8. Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-lwe and security for key dependent messages. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 505–524. Springer, 2011.
9. Ran Canetti and Susan Hohenberger. Chosen-ciphertext secure proxy re-encryption. In Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson, editors, *Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, Virginia, USA, October 28-31, 2007*, pages 185–194. ACM, 2007.
10. Nishanth Chandran, Melissa Chase, Feng-Hao Liu, Ryo Nishimaki, and Keita Xagawa. Re-encryption, functional re-encryption, and multi-hop re-encryption: A framework for achieving obfuscation-based security and instantiations from lattices. In Hugo Krawczyk, editor, *Public-Key Cryptography - PKC 2014 - 17th International Conference on Practice and Theory in Public-Key Cryptography, Buenos Aires, Argentina, March 26-28, 2014. Proceedings*, volume 8383 of *Lecture Notes in Computer Science*, pages 95–112. Springer, 2014.
11. Nishanth Chandran, Melissa Chase, and Vinod Vaikuntanathan. Functional re-encryption and collusion-resistant obfuscation. In Ronald Cramer, editor, *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, volume 7194 of *Lecture Notes in Computer Science*, pages 404–421. Springer, 2012.
12. Aloni Cohen. What about bob? the inadequacy of CPA security for proxy re-encryption. *IACR Cryptology ePrint Archive*, 2017:785, 2017.
13. Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. On post-compromise security. In *IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016*, pages 164–178. IEEE Computer Society, 2016.

14. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 559–585. Springer, 2016.
15. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 324–348, 2017.
16. Adam Everspaugh, Kenneth G. Paterson, Thomas Ristenpart, and Samuel Scott. Key rotation for authenticated encryption. *IACR Cryptology ePrint Archive*, 2017:527, 2017.
17. Georg Fuchsbauer, Chethan Kamath, Karen Klein, and Krzysztof Pietrzak. Adaptively secure proxy re-encryption. *IACR Cryptology ePrint Archive*, 2018:426, 2018.
18. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 169–178. ACM, 2009.
19. Susan Hohenberger, Guy N. Rothblum, Abhi Shelat, and Vinod Vaikuntanathan. Securely obfuscating re-encryption. In Salil P. Vadhan, editor, *Theory of Cryptography, 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21-24, 2007, Proceedings*, volume 4392 of *Lecture Notes in Computer Science*, pages 233–252. Springer, 2007.
20. Zahra Jafargholi, Chethan Kamath, Karen Klein, Ilan Komargodski, Krzysztof Pietrzak, and Daniel Wichs. Be adaptive, avoid overcommitting. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 133–163. Springer, 2017.
21. Ela Lee. Improved security notions for proxy re-encryption to enforce access control. *Cryptology ePrint Archive*, Report 2017/824, 2017. <http://eprint.iacr.org/2017/824>.
22. Anja Lehmann and Björn Tackmann. Updatable encryption with post-compromise security. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 685–716. Springer, 2018.
23. Benoît Libert and Damien Vergnaud. Unidirectional chosen-ciphertext secure proxy re-encryption. In Ronald Cramer, editor, *Public Key Cryptography - PKC 2008, 11th International Workshop on Practice and Theory in Public-Key Cryptography, Barcelona, Spain, March 9-12, 2008. Proceedings*, volume 4939 of *Lecture Notes in Computer Science*, pages 360–379. Springer, 2008.
24. Adriana Lopez-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. *Cryptology ePrint Archive*, Report 2013/094, 2013. <https://eprint.iacr.org/2013/094>.

25. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.
26. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
27. Steven Myers and Adam Shull. Efficient hybrid proxy re-encryption for practical revocation and key rotation. *IACR Cryptology ePrint Archive*, 2017:833, 2017.
28. OWASP. Cryptographic storage cheat sheet. [https://www.owasp.org/index.php/Cryptographic\\_Storage\\_Cheat\\_Sheet](https://www.owasp.org/index.php/Cryptographic_Storage_Cheat_Sheet), 2018. Accessed 9th October 2018.
29. PCI Security Standards Council 2018. Payment Card Industry (PCI) data security standard (version 3.2.1). 2018.
30. Yuriy Polyakov, Kurt Rohloff, Gyana Sahu, and Vinod Vaikuntanathan. Fast proxy re-encryption for publish/subscribe systems. *ACM Trans. Priv. Secur.*, 20(4):14:1–14:31, 2017.

## A Definitions for symmetric PRE

Here we give definitions for the symmetric setting. We leave out some security games where the change is analogous, but explicitly give those which contains lists to clarify how those lists are updated.

**Definition 9.** A symmetric Proxy Re-Encryption (symPRE) scheme consists of the following algorithms:

- $\text{sym.Setup}(1^\lambda) \rightarrow \text{params}$ : Outputs a set of public parameters, including the message space and ciphertext space. Note that  $\text{params}$  is input to every subsequent algorithm, but we leave it out for compactness of notation. We often omit the  $\text{Setup}$  algorithm for the same reason.
- $\text{sym.KeyGen}(1^\lambda) \rightarrow k$ : Generates a secret key
- $\text{sym.Enc}(k, m, \ell) \rightarrow C$ : Encrypts a message  $m$  using a key  $k$ , producing a ciphertext at level  $\ell$ . We often leave out  $\ell$  for compactness.
- $\text{sym.Dec}(k, C) \rightarrow m' \cup \perp$ : Decrypts a ciphertext  $C$  to produce either an element of the message space  $m'$  or the error symbol
- $\text{sym.ReKeyGen}(k_i, k_j) \rightarrow \Delta_{i,j} \cup \perp$ : Takes current key  $k_i$  and next key  $k_j$  and outputs an update token  $\Delta_{i,j}$ , or  $\perp$  when  $i = j$ . This last condition is often left out of constructions for compactness.
- $\text{sym.ReEnc}(\Delta_{i,j}, C) \rightarrow C'$ : Takes a ciphertext  $C$  under  $k_i$  and outputs a new ciphertext  $C'$  under  $k_j$ .

A symmetric PRE scheme is correct if for all  $m \in \mathcal{M}$ ,  $k \xleftarrow{\$} \text{sym.KeyGen}(1^\lambda)$ :

$$\text{sym.Dec}(k, \text{Enc}(k, m)) \rightarrow m$$

and if for all  $C \in \mathcal{C}$  such that  $\text{sym.Dec}(k_i, C) \rightarrow m$ :

$$\text{sym.Dec}(k_j, \text{sym.ReEnc}(\Delta_{i,j}, C)) \rightarrow m$$

where  $k_i, k_j, \xleftarrow{\$} \text{sym.KeyGen}(1^\lambda)$  and  $\Delta_{i,j} \leftarrow \text{sym.ReKeyGen}(k_i, k_j)$ .

$\text{Sym-CPA}_{\mathcal{A}}^{b, \mathcal{SE}}(1^\lambda)$	$\mathcal{O}_{\text{KeyGen}}(1^\lambda)$	$\mathcal{O}_{\text{challenge}}(i, m_0, m_1)$
$\kappa = 0$	$\kappa = \kappa + 1$	<b>if OR</b> $ m_0  \neq  m_1 $ : <b>return</b> $\perp$
$b' \leftarrow \mathcal{A}_1^{\mathcal{O}_{\text{KeyGen}}, \mathcal{O}_{\text{challenge}}}(1^\lambda)$	$k_\kappa \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	$C \xleftarrow{\$} \text{Enc}(k_\kappa, m_b)$
<b>return</b> $b' = b$		<b>return</b> $C$

Figure 10: The Sym-CPA, the symmetric variant of the PKE-CPA game.

**Definition 10.** An encryption scheme  $\mathcal{SE}$  is  $\epsilon$ -Indistinguishable against Chosen Plaintext Attacks ( $\epsilon$ -sym-CPA-secure) if for all PPT adversaries  $\mathcal{A}$ :

$$\left| \Pr \left[ \text{Sym-CPA}_{\mathcal{A}}^{0, \mathcal{SE}}(1^\lambda) = 1 \right] - \Pr \left[ \text{Sym-CPA}_{\mathcal{A}}^{1, \mathcal{SE}}(1^\lambda) = 1 \right] \right| \leq \epsilon,$$

where Sym-CPA is defined in Figure 10. If  $\epsilon$  is negligible as parameterised by the security parameter  $\lambda$ , then we say the scheme is Indistinguishable against Chosen Plaintext Attacks (sym-CPA-secure).

A symmetric PRE scheme  $\text{symPRE}$  is ( $\epsilon$ -sym-CPA-secure) if the encryption scheme given by  $\text{symPKE} = \{\text{symPRE.KeyGen}, \text{symPRE.Enc}, \text{symPRE.Dec}\}$  is  $\epsilon$ -Sym-CPA secure.

$\text{SH}_{\mathcal{A}}^{b, \mathcal{PRE}}(1^\lambda, 1^\kappa)$	$\mathcal{O}_{\text{challenge}}((i, j), m^*, \ell^*)$
$(\text{pk}_\iota, \text{sk}_\iota)_{\iota \in [\kappa]} \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	<b>if</b> $\ell^* > L - 1$ : <b>return</b> $\perp$
$\{\Delta_{i,j} \xleftarrow{\$} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)\}_{(i,j) \in [\kappa]^2}$	$C \leftarrow \text{Enc}(\text{pk}_i, m^*, \ell^*)$
$b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{challenge}}}(1^\lambda, (\text{pk}_\iota, \text{sk}_\iota)_{\iota \in [\kappa]}, \{\Delta_{i,j}\}_{(i,j) \in [\kappa]^2})$	<b>if</b> $b = 0$ :
<b>return</b> $b' = b$	$C^{(b)} \xleftarrow{\$} \text{ReEnc}(\Delta_{i,j}, C)$
	<b>elseif</b> $b = 1$ :
	$C^{(b)} \xleftarrow{\$} \text{Enc}(\text{pk}_j, m^*, \ell^* + 1)$
	<b>return</b> $(C, C^{(b)})$

Figure 11: Experiments for the symmetric source-hiding property.  $L$  is the number of times a ciphertext can be re-encrypted without breaking the correctness conditions.

**Definition 11.** A symmetric PRE scheme  $f\mathcal{PRE}$  is said to be  $\epsilon$ -post-compromise secure ( $\epsilon$ -PCS) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$\left| \Pr \left[ \text{sSymPostComp}_{\mathcal{A}}^{0, f\mathcal{PRE}}(1^\lambda) = 1 \right] - \Pr \left[ \text{sSymPostComp}_{\mathcal{A}}^{1, f\mathcal{PRE}}(1^\lambda) = 1 \right] \right| \leq \epsilon,$$

where  $\text{sSymPostComp}_{\mathcal{A}}^{b, f\mathcal{PRE}}$  is defined in Figure 13.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme is post-compromise secure (PCS).



$\mathcal{O}_{\text{KeyGen}}(1^\lambda)$	$\mathcal{O}_{\text{Corrupt}}(i)$	$\mathcal{O}_{\text{Enc}}(i, m)$	$\mathcal{O}_{\text{ReKeyGen}}(i, j)$
$\kappa = \kappa + 1$	$\mathcal{K}_{\text{corrupted}}.\text{add}(k_i)$	$C \stackrel{\$}{\leftarrow} \text{Enc}(k_i, m)$	$\Delta_{i,j} \stackrel{\$}{\leftarrow} \text{ReKeyGen}(k_i, k_j)$
$k_{\kappa} \stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda)$	<b>return</b> $k_i$	$\mathcal{C}_{\text{honest}}.\text{add}(i, C)$	$\mathcal{T}_{\text{honest}}.\text{add}(i, j, \Delta_{i,j})$
$\mathcal{DRG}.\text{add}(v_\kappa)$		$\mathcal{C}_{\text{msg}}[(i, C)] = m$	$\mathcal{DRG}.\text{add}(\vec{e}_{i,j})$
		<b>return</b> $C$	<b>return</b> $\Delta_{i,j}$

Figure 12: Common oracles used in security games for symmetric PRE.

$\text{sSymPostComp}_{\mathcal{A}}^{b, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda)$	$\mathcal{O}_{\text{ReEnc}}(C, i, j, [\Delta_{i,j}])$
$\mathcal{K}_{\text{chal}}, \mathcal{K}_{\text{corrupted}}, \mathcal{C}_{\text{honest}}, \mathcal{C}_{\text{chal}}, \mathcal{C}_{\text{msg}}, \mathcal{T}_{\text{honest}}, \mathcal{DRG} = \emptyset$	<b>if</b> $\Delta_{i,j}$ <b>given</b> :
$\kappa = 0, \text{called} = \text{false}$	<b>if</b> $(i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}}$ : <b>return</b> $\perp$
$\text{state} \leftarrow \mathcal{A}_0^{\mathcal{O}_{\text{KeyGen}}, \mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{ReEnc}}, \mathcal{O}_{\text{Corrupt}}}(1^\lambda)$	<b>else</b> : $\Delta_{i,j} \leftarrow \text{ReKeyGen}(k_i, k_j)$
$b' \leftarrow \mathcal{A}_1^{\mathcal{O}_{\text{KeyGen}}, \mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{ReKeyGen}}, \mathcal{O}_{\text{ReEnc}}, \mathcal{O}_{\text{challenge}}}(1^\lambda, \text{state})$	<b>if</b> $(i, C) \notin \mathcal{C}_{\text{honest}}$ : <b>return</b> $\perp$
$\mathcal{K}_{\text{chal}} \leftarrow \text{UpdateChallengeKeys}(\mathcal{K}_{\text{chal}}, \mathcal{DRG})$	$C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C)$
<b>if</b> $\mathcal{K}_{\text{chal}} \cap \mathcal{K}_{\text{corrupted}} \neq \emptyset$ : <b>return</b> $\perp$	$\mathcal{C}_{\text{honest}}.\text{add}(j, C')$
<b>return</b> $b' = b$	$\mathcal{C}_{\text{msg}}[(j, C')] = \mathcal{C}_{\text{msg}}[(i, C)]$
	<b>if</b> $(i, C) \in \mathcal{C}_{\text{chal}}$ :
	$\mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(k_j)$
	<b>return</b> $C'$
$\mathcal{O}_{\text{challenge}}(C_0, C_1, i, j, \Delta_{i,j})$	
<b>if</b> $ C_0  \neq  C_1 $ <b>OR</b> <b>called</b> = true : <b>return</b> $\perp$	
<b>if</b> $(i, C_0), (i, C_1) \notin \mathcal{C}_{\text{honest}}$ <b>OR</b> $(i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}}$ : <b>return</b> $\perp$	
$C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C_b)$	
$\mathcal{C}_{\text{msg}}[(j, C')] = \mathcal{C}_{\text{msg}}[(i, C_b)]$	
$\mathcal{C}_{\text{honest}}.\text{add}(j, C'), \mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(k_j)$	
<b>called</b> $\leftarrow$ true	
<b>return</b> $C'$	

Figure 13: The symmetric variant of PostComp game.

We give the equivalent lemma to Lemma 5 for symmetric PRE.

**Lemma 11.**  $\text{IND-UPD} \not\Rightarrow \text{symmetric PCS}$ .

*Proof.* We present the following counterexample for symmetric PRE, showing that  $\text{IND-UPD} \not\Rightarrow \text{symmetric PCS}$ . RISE [22] is proven to be IND-UPD. Here we adapt RISE for general proxy re-encryption as this fits more with our notation, but we observe that the original construction defined as an updatable encryption scheme is also sufficient for the proof.

- $\text{RISE}'.\text{KeyGen}(1^\lambda) : x \stackrel{\$}{\leftarrow} \mathbb{Z}_q^*, (\text{pk}, \text{sk}) = (x, g^x)$
- $\text{RISE}'.\text{Enc}(\text{pk}, m) : r \stackrel{\$}{\leftarrow} \mathbb{Z}_q^*, C \leftarrow (\text{pk}^r, g^r \cdot m)$

- $\text{RISE}'.\text{Dec}(\text{sk}, m) : m' \leftarrow C_1 \cdot C_0^{-1/\text{sk}}$
- $\text{RISE}'.\text{ReKeyGen}(\text{sk}_i, \text{sk}_j) : \Delta_{i,j} = (\text{sk}_j/\text{sk}_i, \text{pk}_j) = (x_j/x_i, g^{x_j})$
- $\text{RISE}'.\text{ReEnc}(\Delta_{i,j}, C) : (\Delta, y') = \Delta_{i,j}, r' \xleftarrow{\$} \mathbb{Z}_q^*, C' \leftarrow (C_0^\Delta \cdot y'^{r'}, C_1 \cdot g^{r'})$

Given  $\Delta_{i,j} = (x_j/x_i, g^{x_j})$ , and  $x_i, \mathcal{A}$  can compute  $x_j$  and use this to decrypt the challenge ciphertext. This is because RISE is not unidirectional it is not compromise secure by Lemma 2.

## B Asymmetric IND-UPD

IND-UPD [22, definition 3] is a post-compromise security notion for updatable encryption schemes, which are a variant of symmetric PCS schemes. Here we adapt IND-UPD for the public-key setting.

**Definition 12.** A PRE scheme  $\mathcal{PRE}$  is said to be (selectively)  $\epsilon$ -pkIND-UPD-secure if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$|\Pr[\text{pkIND-UPD}_{\mathcal{A}}^0(1^\lambda) = 1] - \Pr[\text{pkIND-UPD}_{\mathcal{A}}^1(1^\lambda) = 1]| \leq \epsilon,$$

where  $\text{pkIND-UPD}_{\mathcal{A}}^{b, \mathcal{PRE}}$  is defined in Figure 14.

If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme is (selectively) pkIND-UPD-secure.

## C BV-PRE

Here we give the BV-PRE scheme [30] as well as the full proof of Theorem 4. This scheme is not fully public-key as it uses additional ‘public’ key ‘ $\text{pk}'_B$ ’ for the target key, which cannot be made public without exposing  $\text{sk}_A$  to anyone with the update token.

**Key Resampling.** The key resampling technique [8], can be found in Figure 16. We use this in our construction in Figure 8 to obtain PCS.

### C.1 Correctness of our modified BV-PRE scheme

Much of the analysis used to argue correctness from [30] holds for our modified scheme in Figure 8. Therefore, we keep the discussion fairly brief and refer to [30] for full details.

Recall from [30] that on input ciphertext  $(c_0, c_1)$  and secret key  $s$ , the decryption algorithm computes  $c_0 - c_1s$  and then performs a reduction modulo  $p$ . In the case that  $(c_0, c_1)$  is a ciphertext produced by the Enc algorithm, we have that  $c_0 - c_1s = p(ev + e_0 + e_1s) + m \bmod q$  and that decryption is successful when  $p(ev + e_0 + e_1s) + m$  (i.e. the error plus the message) does not wrap around modulo  $q$ . However, the correctness condition of decryption is different when considering ciphertexts output by the ReEnc algorithm. To see how, let

$\text{pkIND-UPD}_{\mathcal{A}}^{b, \mathcal{PRE}}(1^\lambda)$ <hr/> $\mathcal{K}_{\text{corrupted}}, \mathcal{C}_{\text{honest}}, \mathcal{K}_{\text{chal}}, \mathcal{DRG} = \emptyset$ $\kappa = 0, c = \perp$ $\text{state} \leftarrow \mathcal{A}_0^{\text{O}_{\text{KeyGen}}^{\text{pkIU}}, \text{O}_{\text{Corrupt}}, \text{O}_{\text{Enc}}, \text{O}_{\text{LearnTok}}, \text{O}_{\text{ReEnc}}^{\text{pkIU}}}(1^\lambda)$ $b' \leftarrow \mathcal{A}_1^{\text{O}_{\text{KeyGen}}^{\text{pkIU}}, \text{O}_{\text{Enc}}, \text{O}_{\text{LearnChal}}, \text{O}_{\text{ReEnc}}, \text{O}_{\text{challenge}}^{\text{pkIU}}}(1^\lambda, \text{state})$ $\mathcal{K}_{\text{chal}} \leftarrow \text{UpdateChallengeKeys}(\mathcal{K}_{\text{chal}}, \mathcal{DRG})$ $\text{if } \mathcal{K}_{\text{chal}} \cap \mathcal{K}_{\text{corrupted}} \neq \emptyset : \text{return } 0$ $\text{return } b'$	$\text{O}_{\text{KeyGen}}^{\text{pkIU}}(1^\lambda)$ <hr/> $\kappa = \kappa + 1$ $\mathcal{DRG}.\text{add}(v_\kappa)$ $(\text{pk}_\kappa, \text{sk}_\kappa) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ $\text{if } \kappa > 1 :$ $\quad \Delta_{\kappa-1, \kappa} \leftarrow \text{ReKeyGen}(\text{sk}_{\kappa-1}, \text{pk}_\kappa)$ $\text{if } c \neq \perp :$ $\quad C_\kappa^* \leftarrow \text{ReEnc}(\Delta_{\kappa-1, \kappa}, C_{\kappa-1}^*)$ $\text{return } \text{pk}_\kappa$
$\text{O}_{\text{LearnTok}}^{\text{pkIU}}(i)$ <hr/> $\text{if } i = c : \text{return } \perp$ $\text{if } \mathcal{PRE} \text{ is unidirectional} :$ $\quad \mathcal{DRG}.\text{add}(\vec{e}_{i-1, i})$ $\text{elseif } \mathcal{PRE} \text{ is bidirectional} :$ $\quad \mathcal{DRG}.\text{add}(e_{i, i-1}(\text{undirected edge}))$ $\text{return } \Delta_{i-1, i}$	$\text{O}_{\text{LearnChal}}^{\text{pkIU}}()$ <hr/> $\mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_\kappa)$ $\text{return } C_\kappa^*$
$\text{O}_{\text{ReEnc}}^{\text{pkIU}}(C)$ <hr/> $\text{if } (\kappa - 1, C) \notin \mathcal{C}_{\text{honest}} : \text{return } \perp$ $\text{if } \text{ReEnc} \text{ is deterministic} :$ $\quad \text{if } C = \bar{C}_0 \text{ OR } C = \bar{C}_1 : \text{return } \perp$ $C' \xleftarrow{\$} \text{ReEnc}(\Delta_{\kappa-1, \kappa}, C)$ $\mathcal{C}_{\text{honest}}.\text{add}(\kappa, C')$ $\text{return } C'$	$\text{O}_{\text{challenge}}^{\text{pkIU}}(C_0, C_1)$ <hr/> $\text{if } c \neq \perp \text{ OR }  C_0  \neq  C_1  : \text{return } \perp$ $\text{if } (\kappa, C_0), (\kappa, C_1) \notin \mathcal{C}_{\text{honest}} : \text{return } \perp$ $\kappa = \kappa + 1, \mathcal{DRG}.\text{add}(v_\kappa)$ $c = \kappa, \bar{C}_0 = C_0, \bar{C}_1 = C_1$ $(\text{pk}_\kappa, \text{sk}_\kappa) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ $\Delta_{\kappa-1, \kappa} \leftarrow \text{ReKeyGen}(\text{sk}_{\kappa-1}, \text{pk}_\kappa)$ $C_\kappa^* \leftarrow \text{ReEnc}(\Delta_{\kappa-1, \kappa}, C_b)$ $\mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_\kappa)$ $\text{return } (\text{pk}_\kappa, C_\kappa^*)$

Figure 14: The pkIND-UPD game, based on IND-UPD [22] adapted to the public-key setting.

$(c'_0, c'_1) \xleftarrow{\$} \text{ReEnc}(\Delta_{A \rightarrow B}, (c_0, c_1))$ . Using notation consistent with Figure 16, we have that

$$\begin{aligned}
c'_0 - c'_1 s_B &= c_0 + \sum_{i=0}^{\lfloor \log(q)/r \rfloor} c_1^{(i)} (\theta_i - s_A (2^r)^i - s_B \beta_i) + (\theta^{\text{proxy}} - \beta^{\text{proxy}} s_B) \\
&= c_0 - c_1 s_A + \sum_{i=0}^{\lfloor \log(q)/r \rfloor} c_1^{(i)} (\theta_i - s_B \beta_i) + (\theta^{\text{proxy}} - \beta^{\text{proxy}} s_B) \\
&= c_0 - c_1 s_A + \sum_{i=0}^{\lfloor \log(q)/r \rfloor} c_1^{(i)} \cdot p(ev_i + e'_i s_B + e''_i) + p(e^{\text{proxy}} v^{\text{proxy}} + e'^{\text{proxy}} s_B + e''^{\text{proxy}}),
\end{aligned}$$

Setup( $1^\lambda$ )	KeyGen( $1^\lambda$ )	Enc(pk, m)
Choose positive integers $q, n$	$a \xleftarrow{\$} \mathcal{U}_q$	$(a, b) \leftarrow \text{pk}$
Choose plaintext modulus $p$	$s, e \xleftarrow{\$} \chi_e$	$v, e_0, e_1 \xleftarrow{\$} \chi_e$
Choose key switching window $r$	$b = a \cdot s + pe$	$c_0 = b \cdot v + pe_0 + m$
<b>return</b> $pp = (q, n, p, r)$	$\text{sk} = s, \text{pk} = (a, b)$	$c_1 = a \cdot v + pe_1$
	<b>return</b> (sk, pk)	<b>return</b> $c = (c_0, c_1)$
	<b>Preprocess</b> ( $1^\lambda, \text{sk}_B$ )	
<b>Dec</b> (sk, c)	$s_b \leftarrow \text{sk}_b$	
$s \leftarrow \text{sk}, (c_0, c_1) \leftarrow c$	<b>for</b> $i \in \{0, 1, \dots, \lceil \log_2(q)/r \rceil\}$ :	
$m' = c_0 - s \cdot c_1 \pmod p$	$\beta_i \xleftarrow{\$} \mathcal{U}_q$	
<b>return</b> $m'$	$e_i \xleftarrow{\$} \chi_e$	
	$\theta_i = \beta_i \cdot s_B + pe_i$	
	<b>return</b> $\text{'pk'}'_B = \{(\beta_i, \theta_i)\}_{i=0}^{\lceil \log_2(q)/r \rceil}$	
	<b>ReEnc</b> ( $\Delta_{A,B}, c$ )	
<b>ReKeyGen</b> ( $\text{sk}_A, \text{'pk'}'_B$ )	$\{(\beta_i, \gamma_i)\}_{i=0}^{\lceil \log_2(q)/r \rceil} \leftarrow \Delta_{A,B}$	
$s_A \leftarrow \text{sk}_A, \{(\beta_i, \theta_i)\}_{i=0}^{\lceil \log_2(q)/r \rceil} \leftarrow \text{'pk'}'_B$	$(c_0, c_1) \leftarrow c$	
<b>for</b> $i \in \{0, 1, \dots, \lceil \log_2(q)/r \rceil\}$ :	$c'_0 = c_0 + \sum_{i=0}^{\lceil \log_2(q)/r \rceil} (c_1^{(i)} \cdot \gamma_i)$	
$\gamma_i = \theta_i - s_A \cdot (2^r)^i$	$c'_1 = \sum_{i=0}^{\lceil \log_2(q)/r \rceil} (c_1^{(i)} \cdot \beta_i)$	
$\Delta_{A \rightarrow B} = \{(\beta_i, \gamma_i)\}_{i=0}^{\lceil \log_2(q)/r \rceil}$	<b>return</b> $c' = (c'_0, c'_1)$	
<b>return</b> $\Delta_{A,B}$		

Figure 15: BV-PRE [30] where  $\mathcal{U}_q$  is the discrete uniform distribution over  $R_q$  and  $\chi_e$  is a  $B_e$ -bounded discrete Gaussian error distribution.

where all arithmetic is done over the integers modulo  $q$ . This shows that the error grows by an additive term of  $pE$  on each re-encryption where

$$E = \sum_{i=0}^{\lceil \log(q)/r \rceil} c_1^{(i)} \cdot (ev_i + e'_i s_B + e''_i) + (e^{\text{proxy}} v^{\text{proxy}} + e'^{\text{proxy}} s_B + e''^{\text{proxy}}).$$

Assume that  $\chi_e$  is  $(B, \delta)$ -bounded for some small  $\delta$  (we quantify these values later). Then we have that the noise grows by at most

$$p\|E\|_\infty \leq pn(2^r - 1)(\lceil \log(q)/r \rceil + 1)(2B^2n + B) + p(2B^2n + B) =: G(n, q, r, B) \quad (2)$$

on each re-encryption taking into consideration the multiplication of degree  $n - 1$  polynomials. Therefore, after  $\ell$  re-encryptions, the noise has grown by an additive term of size  $\ell G$ . Therefore, the condition for successful decryption after  $\ell$  re-encryptions is

$$p(2B^n + B) + \ell G(n, q, r, B) + p/2 \leq q/2 \quad (3)$$

ReSample(pk)
$(a, b) \leftarrow \text{pk}$
$v, e' \xleftarrow{\$} \chi_e$
$e'' \xleftarrow{\$} \chi_e$
$\bar{a} = av + pe'$
$\bar{b} = bv + pe''$
<b>return</b> $\text{pk} = (\bar{a}, \bar{b})$

Figure 16: Key resampling technique [8].

where  $G$  is define in Equation (2). Note that if our error distribution has  $\sigma > \omega(\log n)$ , then we can set  $B = \sigma\sqrt{n}$  and  $\delta = 2^{-n+1}$  [24, 26]. In order to choose parameters of the system, one needs to ensure that Equation 3 is satisfied. Note that this analysis is a worst-case one. It is shown in [30] that the central limit theorem can be used to essentially alter the form of  $B$  and change all occurrences of  $n$  into occurrences of  $\sqrt{n}$  in Equation (3) by allowing for a tunable failure probability. We omit this more practical method of correctness analysis for brevity.

## C.2 Full proof of Theorem 4

*Proof.* Let  $N$  be the number of honest parties and  $M$  be the total number of parties where  $M$  is polynomially bounded. First of all, we label the nodes in our  $\mathcal{DRG}$  using a topological ordering (from sinks to sources) such that there are no edges from  $i$  to  $k$  if  $i < k$ . In doing so, we have that the honest and corrupt entities have labels in the set  $\Gamma_H = \{1, \dots, N\}$  and  $\Gamma_C = \{N+1, \dots, M\}$  respectively. For  $i = 1, \dots, M$ , we denote the secret and public keys of node  $i$  as

$$\text{sk}_i = s_i, \quad \text{pk}_i = (a_i, b_i = a_i \cdot s_i + e_i) \quad (4)$$

where  $a_i \xleftarrow{\$} R_q$  and  $s_i, e_i \xleftarrow{\$} \chi_e$ . Also, for  $j = 1, \dots, Q_i$  where  $Q_i$  is a polynomial bounding the number of queries to  $\text{ReSample}(\text{pk}_i)$ , we denote the output of the  $j^{\text{th}}$  call of the form  $\text{ReSample}(\text{pk}_i)$  as the pair

$$a'_{i,j} = a_i v_{i,j} + pe'_{i,j}, \quad b'_{i,j} = b_i v_{i,j} + pe''_{i,j} \quad (5)$$

where  $v_{i,j}, e'_{i,j}, e''_{i,j} \xleftarrow{\$} \chi_e$ . We now state our main hybrid games between the adversary  $\mathcal{A}$  and challenger  $\mathcal{C}$ :

- **Game<sub>0</sub>**: Run the  $\text{PostComp}^{b, \mathcal{PR}\mathcal{E}}$  security game where the bit  $b \xleftarrow{\$} \{0, 1\}$  uniformly. Denoting the adversary  $\mathcal{A}$ 's output in the  $\text{PostComp}^{b, \mathcal{PR}\mathcal{E}}$  as  $b'$ , the output of **Game<sub>0</sub>** is 1 if  $b' = b$  and 0 otherwise.
- **Game<sub>i</sub>** for  $i = 1, \dots, N$ : Same as **Game<sub>i-1</sub>** except that the challenger  $\mathcal{C}$  replaces  $\text{sk}_i, \text{pk}_i$  and all terms  $a'_{i,\cdot}, b'_{i,\cdot}$  in (5) by random uniform values.

- $\text{Game}_{N+1}$ :  $\mathcal{C}$  behaves the same as  $\text{Game}_N$  except that all challenge ciphertexts output by the challenge oracle are replaced by uniform random values.

Note that  $\text{Game}_0$  challenges an adversary to distinguish between the cases when it plays either  $\text{PostComp}^{0, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda)$  and  $\text{PostComp}^{1, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda)$ . Therefore, we aim to show that there is no PPT adversary that wins  $\text{Game}_0$  with probability a non-negligible amount more than  $1/2$ . The adversary information theoretically has no advantage in  $\text{Game}_{N+1}$  i.e.

$$\Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{N+1}(\mathcal{A}) \right] - 1/2 = 0.$$

This means that we can prove security by showing that  $\text{Game}_0$  is computationally indistinguishable from  $\text{Game}_{N+1}$  according to the RLWE assumption. We do so by proving the following two lemmas:

**Lemma 12.**  *$\text{Game}_{i-1}$  is computationally indistinguishable from  $\text{Game}_i$  for  $i = 1, \dots, N$  under the RLWE assumption. In particular, let  $Q_i$  be the number of times  $\text{ReSample}(\text{pk}_i)$  is called during the  $\text{PostComp}$  game. Then for any PPT adversary  $\mathcal{A}$ , there exists a RLWE distinguisher  $\mathcal{D}$  such that*

$$\left| \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{i-1}^{k-1}(\mathcal{A}) \right] - \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{i-1}^k(\mathcal{A}) \right] \right| \leq (1+Q_i) \mathbf{Adv}_{\text{RLWE}_{\phi, q, \chi_e}}(\mathcal{D}).$$

*Proof.* Let  $Q_i$  be the polynomially bounded number of times that  $\text{ReSample}(\text{pk}_i)$  is called. We define a sequence of sub-hybrids between  $\text{Game}_{i-1}$  and  $\text{Game}_i$ :

- $\text{Game}_{i-1}^0$ : The same as  $\text{Game}_{i-1}$  apart from the fact that the challenger replaces the public key  $\text{pk}_i$  with a uniform random value.
- $\text{Game}_{i-1}^k$  for  $k = 1, \dots, P_i$ : The same as  $\text{Game}_{i-1}^{k-1}$  apart from that the  $k^{\text{th}}$  call to  $\text{ReSample}(\text{pk}_i)$  is replaced by a uniform random value.

Firstly, since node  $i$  is not corrupted by the  $\text{PostComp}$  adversary  $\mathcal{A}$ , a RLWE adversary  $\mathcal{D}$  can embed a RLWE challenge into  $\text{pk}_i$  to simulate  $\text{Game}_{i-1}$  or  $\text{Game}_{i-1}^0$  depending on whether its input challenge was uniform or not. In particular, since all calls to the resample algorithm are uniform for public keys associated to nodes with label  $\leq i-1$ , the  $\text{ReKeyGen}$  oracle outputs uniform random values when  $\mathcal{A}$  asks for a re-encryption key with source node  $i$ . Therefore,  $\mathcal{D}$  simply returns uniform values on these particular  $\text{ReKeyGen}$  queries. If  $\text{Game}_{i-1}$  and  $\text{Game}_{i-1}^0$  are distinguishable, then  $\mathcal{D}$  can solve the RLWE problem. Therefore,

$$\left| \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{i-1}(\mathcal{A}) \right] - \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{i-1}^0(\mathcal{A}) \right] \right| \leq \mathbf{Adv}_{\text{RLWE}_{\phi, q, \chi_e}}(\mathcal{D}).$$

Secondly, a RLWE adversary  $\mathcal{D}$  can simulate either  $\text{Game}_{i-1}^{k-1}$  or  $\text{Game}_{i-1}^k$  by performing the following:

1. Query the RLWE oracle twice obtain  $(a, d_1), (b, d_2)$  and set  $\text{pk}_i = (a, b)$ .
2. Answer the  $k^{\text{th}}$  instance of the query  $\text{ReSample}(\text{pk}_i)$  as the pair  $(d_1, d_2)$
3. Answer all other queries as in  $\text{Game}_{i-1}^{k-1}$

Once again, if  $\text{Game}_{i-1}^{k-1}$  and  $\text{Game}_{i-1}^k$  are distinguishable, then  $\mathcal{D}$  can solve the RLWE problem. Therefore,

$$\left| \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{i-1}^{k-1}(\mathcal{A}) \right] - \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{i-1}^k(\mathcal{A}) \right] \right| \leq \mathbf{AdvRLWE}_{\phi, q, \chi_e}(\mathcal{D}).$$

Using the triangle inequality implies that

$$\left| \Pr \left[ \text{PostComp}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda) = 1 \right] - \frac{1}{2} \right| \leq \epsilon (\rho \cdot Q_{rk} + Q_{re} + N) \cdot \mathbf{AdvRLWE}_{\phi, q, \chi_e}(\mathcal{D}).$$

and applying the RLWE assumption completes the proof of this lemma.  $\square$

**Lemma 13.** *For any adversary  $\mathcal{A}$ ,  $\Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_N(\mathcal{A}) \right] = \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_{N+1}(\mathcal{A}) \right]$ .*

*Proof.* In both  $\text{Game}_N$  and  $\text{Game}_{N+1}$ , the adversary submits two ciphertexts  $C_0 = (c_{0,0}, c_{0,1})$  and  $C_1 = (c_{1,0}, c_{1,1})$  when it queries the challenge oracle. In  $\text{Game}_N$ , the challenger responds by setting  $c_0 = c_{0(b)}$  and  $c_1 = c_{1(b)}$  for uniformly chosen  $b \stackrel{\$}{\leftarrow} \{0, 1\}$  and returning the challenge ciphertext given by

$$c'_0 = c_0 + \theta^* + \sum_{i=0}^{\lfloor \log_2(q)/r \rfloor} (c_1^{(i)} \cdot \gamma_i), \quad c'_1 = \beta^* + \sum_{i=0}^{\lfloor \log_2(q)/r \rfloor} (c_1^{(i)} \cdot \beta_i),$$

where  $\sum_{i=0}^{\lfloor \log_2(q)/r \rfloor} 2^{ri} c_1^{(i)} = c_1$ . In particular,  $\beta^*$  and  $\theta^*$  are the result of calling  $\text{ReSample}()$  in  $\text{Game}_N$  on an honest public key and are therefore uniform random values that are used once and never revealed to  $\mathcal{A}$ . Therefore,  $c'_0$  and  $c'_1$  are independent uniform random values in  $\text{Game}_N$  which is the exact case in  $\text{Game}_{N+1}$ . This argument holds for each of  $\mathcal{A}$ 's challenge oracle queries. This concludes the proof of this lemma.  $\square$

After noting that  $\sum_{i=1}^N Q_i = \rho Q_{rk} + Q_{re}$  and using the triangle inequality, we see that for all PPT adversaries  $\mathcal{A}$ :

$$\left| \Pr \left[ 1 \stackrel{\$}{\leftarrow} \text{Game}_0(\mathcal{A}) \right] - \frac{1}{2} \right| \leq \epsilon,$$

for negligible  $\epsilon$  according to the RLWE assumption. This inequality implies that the games  $\text{PostComp}^{0, \mathcal{P}\mathcal{R}\mathcal{E}}$  and  $\text{PostComp}^{1, \mathcal{P}\mathcal{R}\mathcal{E}}$  are  $\epsilon$ -indistinguishable as required. This concludes the proof of the theorem.  $\square$

## D Adaptive HRA security

Here we give an extension of selective IND-HRA security as described in Section 3.3, which allows the adversary to adaptively corrupt keys in response after receiving challenges.

$\text{adIND-HRA}_{\mathcal{A}}^{b, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda)$ <hr style="border: 0.5px solid black;"/> $\mathcal{K}_{\text{corrupted}}, \mathcal{K}_{\text{chal}}, \mathcal{C}_{\text{chal}}, \mathcal{T}_{\text{honest}}, \mathcal{DRG} = \emptyset$ $\kappa = 0, \text{called} \leftarrow \text{false}$ $b \xleftarrow{\$} \{0, 1\}$ $b' \leftarrow \mathcal{A}^{\text{O}_{\text{KeyGen}}, \text{O}_{\text{Enc}}, \text{O}_{\text{Corrupt}}, \text{O}_{\text{ReEnc}}, \text{O}_{\text{ReEnc}}, \text{O}_{\text{challenge}}}(1^\lambda)$ $\mathcal{K}_{\text{chal}} \leftarrow \text{UpdateChallengeKeys}(\mathcal{K}_{\text{chal}}, \mathcal{DRG})$ <b>if</b> $\mathcal{K}_{\text{corrupted}} \cap \mathcal{K}_{\text{chal}} \neq \emptyset$ : <b>return</b> 0 <b>else return</b> $b'$	$\text{O}_{\text{ReEnc}}(C, i, j, [\Delta_{i,j}])$ <hr style="border: 0.5px solid black;"/> <b>if</b> $\Delta_{i,j}$ given <b>AND</b> $(i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}}$ : <b>return</b> $\perp$ <b>if</b> $\Delta_{i,j}$ not given : $\Delta_{i,j} \xleftarrow{\$} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$ $C' \xleftarrow{\$} \text{ReEnc}(\Delta_{i,j}, C)$ <b>if</b> $(i, C) \in \mathcal{C}_{\text{chal}}$ : $\mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add sk}_j$ <b>return</b> $C'$
$\text{O}_{\text{challenge}}(m_0, m_1, i)$ <hr style="border: 0.5px solid black;"/> <b>if</b> $ m_0  \neq  m_1 $ <b>OR</b> $\text{called} = \text{true}$ : <b>return</b> $\perp$ $C \xleftarrow{\$} \text{ReEnc}(\text{pk}_i, m_b)$ $\mathcal{C}_{\text{chal}}.\text{add}(i, C)$ $\mathcal{K}_{\text{chal}}.\text{add sk}_i$ $\text{called} \leftarrow \text{true}$ <b>return</b> $C$	

Figure 17: The adIND-HRA game. Like the HRA model [12], it allows re-encryptions of non-challenge ciphertexts to compromised keys using  $\text{O}_{\text{ReEnc}}$ . It allows adaptive corruption of keys subject to the trivial win condition that no key which a challenge ciphertext has been learned under can be corrupted.

**Definition 13.** A PRE scheme  $\mathcal{P}\mathcal{R}\mathcal{E}$  is said to be  $\epsilon$ -Adaptively Indistinguishable Honest Re-encryption Plaintext Attacks-secure ( $\epsilon$ -adIND-HRA-secure) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$\left| \Pr \left[ \text{adIND-HRA}_{\mathcal{A}}^{0, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda) = 1 \right] - \Pr \left[ \text{adIND-HRA}_{\mathcal{A}}^{1, \mathcal{P}\mathcal{R}\mathcal{E}}(1^\lambda) = 1 \right] \right| \leq \epsilon, \quad (6)$$

where  $\text{adIND-HRA}_{\mathcal{A}}^{\mathcal{P}\mathcal{R}\mathcal{E}}$  is given in Figure 17. If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme is Adaptively Indistinguishable Honest Re-encryption Plaintext Attacks-secure (adIND-HRA-secure).

In the adIND-HRA game (Figure 17),  $\mathcal{A}$  can adaptively corrupt keys subject to the trivial win condition that they cannot have corrupted a key with which a challenge ciphertext can be decrypted. In order to keep track of the status of keys, the lists  $\mathcal{C}_{\text{chal}}, \mathcal{K}_{\text{chal}}$  and  $\mathcal{K}_{\text{corrupted}}$  are used as well as a directed re-encryption graph  $\mathcal{DRG}$  which tracks update token queries.  $\mathcal{C}_{\text{chal}}$  contains outputs of the challenge oracle  $\text{O}_{\text{challenge}}$  as well as outputs of  $\text{O}_{\text{ReEnc}}$  when given a challenge as input.  $\mathcal{K}_{\text{corrupted}}$  contains the outputs of  $\text{O}_{\text{Corrupt}}$  queries.  $\mathcal{DRG}$  consists of nodes  $v_i$  that represent key pairs, and edges  $\vec{e}_{i,j}$  which are added when  $\text{O}_{\text{ReKeyGen}}(i, j)$  is queried. Using update tokens, the adversary can locally re-encrypt challenge



ciphertexts. Therefore, if a challenge ciphertext is and encryption under  $\text{pk}_i$ , and there exists a sequence of tokens going from  $i$  to  $j$ , then both  $\text{sk}_i$  and  $\text{sk}_j$  are considered challenge keys. Represented using the graph  $\mathcal{DRG}$ , if there is a path from  $v_i$  to  $v_j$  and  $\text{sk}_i$  is a challenge key, then so is  $\text{sk}_j$ . At the end of the game,  $\mathcal{DRG}$  is used to update the list of challenge keys. Then the trivial winning condition translates  $\mathcal{K}_{\text{chal}} \cap \mathcal{K}_{\text{corrupted}}$  being empty.

## E Adaptive Post-Compromise Security

Here we give the explicit definition for adaptive PCS.

$\text{ad-PostComp}_{\mathcal{A}}^{b, \mathcal{PRE}}(1^\lambda)$ <hr style="border: 0.5px solid black;"/> $\mathcal{K}_{\text{corrupted}}, \mathcal{K}_{\text{chal}}, \mathcal{C}_{\text{honest}}, \mathcal{C}_{\text{chal}}, \mathcal{T}_{\text{honest}}, \mathcal{DRG} = \emptyset$ $\kappa = 0, \text{called} = \text{false}$ $b' \leftarrow \mathcal{A}^{\text{OKeyGen}, \text{OEnc}, \text{OCorrupt}, \text{OReEnc}^{\text{aPC}}, \text{Ochallenge}}(1^\lambda)$ $\mathcal{K}_{\text{chal}} \leftarrow \text{UpdateChallengeKeys}(\mathcal{K}_{\text{chal}}, \mathcal{DRG})$ $\text{if } \mathcal{K}_{\text{corrupted}} \cap \mathcal{K}_{\text{chal}} \neq \emptyset :$ $\quad \text{return } \perp$ $\text{else return } b'$	$\text{O}_{\text{ReEnc}}^{\text{aPC}}(C, i, j, [\Delta_{i,j}])$ <hr style="border: 0.5px solid black;"/> $\text{if }  C_0  \neq  C_1  \text{ OR called} = \text{true} : \text{return } \perp$ $\text{if } (i, C_0), (i, C_1) \notin \mathcal{C}_{\text{honest}} : \text{return } \perp$ $\text{if } (i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}} : \text{return } \perp$ $\text{if not given} : \Delta_{i,j} \stackrel{\$}{\leftarrow} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$ $C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C)$ $\text{if } (i, C) \in \mathcal{C}_{\text{honest}} : \mathcal{C}_{\text{honest}}.\text{add}(j, C')$ $\text{if } (i, C) \in \mathcal{C}_{\text{chal}} :$ $\quad \mathcal{C}_{\text{chal}}.\text{add}(j, C')$ $\quad \mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_j)$ $\text{return } C'$
$\text{O}_{\text{challenge}}^{\text{aPC}}(C_0, C_1, i, j, \Delta_{i,j})$ <hr style="border: 0.5px solid black;"/> $\text{if }  C_0  \neq  C_1  : \text{return } \perp$ $\text{if } (i, C_0), (i, C_1) \notin \mathcal{C}_{\text{honest}} : \text{return } \perp$ $\text{if } (i, j, \Delta_{i,j}) \notin \mathcal{T}_{\text{honest}} : \text{return } \perp$ $C' \stackrel{\$}{\leftarrow} \text{ReEnc}(\Delta_{i,j}, C_b)$ $\mathcal{C}_{\text{honest}}.\text{add}(j, C'), \mathcal{C}_{\text{chal}}.\text{add}(j, C'), \mathcal{K}_{\text{chal}}.\text{add}(\text{sk}_j)$ $\text{called} \leftarrow \text{false}$ $\text{return } C'$	

Figure 18: The adaptive post compromise game  $\text{PostComp}$ . This is stronger than the selective version as the adversary can choose to corrupt keys as a result of challenge queries, subject to the trivial win condition.

**Definition 14.** A  $\text{PRE}$  scheme  $\mathcal{PRE}$  is said to have  $\epsilon$ -Adaptive Post-Compromise Security ( $\epsilon$ -adPCS) if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ :

$$|\Pr[\text{ad-PostComp}_{\mathcal{A}}^0(1^\lambda) = 1] - \Pr[\text{ad-PostComp}_{\mathcal{A}}^1(1^\lambda) = 1]| \leq \epsilon,$$

where  $\text{ad-PostComp}_{\mathcal{A}}^{b, \mathcal{PR}\mathcal{E}}$  is defined in Figure 18. If  $\epsilon$  is negligible as parameterised by the security parameter, then we say the scheme has Adaptive Post-Compromise Security (adPCS).

## F Weak key privacy of our construction

**Theorem 7.** *The scheme given in Figure 8 has weak key privacy.*

*Proof.* Once again, the proof replaces the public key and  $\text{ReSample}$  outputs associated to identity labels  $i = 1, \dots, \kappa$  by uniform values one by one. This is permissible due to the RLWE assumption. Define the following sequence of games:

- $\text{Game}_0$ : The  $\text{weakKP}_{\mathcal{A}}^{b, \mathcal{PR}\mathcal{E}}(1^\lambda, 1^\kappa)$  game.
- $\text{Game}'_i$  for  $i = 1, \dots, \kappa$ : The same as  $\text{Game}_{i-1}$  apart from the fact that  $\text{pk}_i$  is replaced with a uniform value.
- $\text{Game}_i$  for  $i = 1, \dots, \kappa$ : The same as  $\text{Game}'_i$  except that the output of calls to  $\text{ReKeyGen}(\cdot, \text{pk}_i)$  are replaced by a uniform value.

The transitions or game hops occur in the order  $\text{Game}_0 \rightarrow \text{Game}'_1 \rightarrow \text{Game}_1 \rightarrow \dots \rightarrow \text{Game}'_\kappa \rightarrow \text{Game}_\kappa$ . We first note that the secret keys  $\text{sk}_1, \dots, \text{sk}_\kappa$  are never used in any of the games so we ignore these throughout this proof.  $\text{Game}_{i-1}$  can be seen to be indistinguishable from  $\text{Game}'_i$  by acknowledging that the difference between these games is that a single RLWE public key is replaced by a uniform value. Next we need to argue that  $\text{Game}'_i$  is indistinguishable from  $\text{Game}_i$ . To do so, we consider a sequence of hybrids between  $\text{Game}'_i$  and  $\text{Game}_i$  denoted as  $\text{Game}'_{i,j}$  where  $j = 0 \dots \lceil \log_2(q)/r \rceil$ . Essentially, in  $\text{Game}'_{i,j}$ , the first  $j$  calls of the form  $\text{ReSample}(pk_i)$  are replaced by uniform random values. Recall that for both these games, the public key  $\text{pk}_i = (a_1, a_2)$  is a uniform random value itself. Therefore, the difference between  $\text{Game}'_{i,j}$  and  $\text{Game}'_{i,j+1}$  is that the  $(j+1)^{\text{th}}$  call to  $\text{ReSample}$  denoted by  $(b_1, b_2)$  either has the form  $(a_1v + e', a_2v + e'')$  or takes the form of a uniform random value. Considering the pairs  $(a_1, b_1)$  and  $(a_2, b_2)$ , an adversary cannot distinguish between  $\text{Game}'_{i,j}$  and  $\text{Game}'_{i,j+1}$  with non-negligible advantage according to the RLWE assumption. This holds for  $j = 0, \dots, \lceil \log_2(q)/r \rceil - 1$ . Since we have  $\text{Game}'_{i,0} = \text{Game}'_i$  and  $\text{Game}'_{i, \lceil \log_2(q)/r \rceil} = \text{Game}_i$ , we can conclude that  $\text{Game}'_i$  is indistinguishable from  $\text{Game}_i$  by using the RLWE assumption multiple times. Iterating through  $i = 0$  to  $\kappa$  completes the proof.  $\square$